University of Nebraska Lincoln

Justin (Gus) Hurwitz
Assistant Professor of Law
(402) 472-1255
ghurwitz@unl.edu

Before the
Federal Communications Commission
Washington, D.C. 20554

In the Matters of	)	
	)	
Protecting and Promoting the Open Internet	)	GN Docket No. 14-28
	)	
Framework for Broadband Internet Service	)	GN Docket No. 10-127
	)	

Comments of Justin (Gus) Hurwitz, Assistant Professor of Law,
University of Nebraska College of Law[1]

These comments are respectfully submitted in response to the Commission's May 15, 2014, Notice of Proposed Rulemaking in GN Docket Number 14-28, relating to the Open Internet proceedings, and related matters. Also included with these comments are copies of recent short writings relating to the issues presented in the Notice.

To summarize, these comments argue that the Commission should not adopt comprehensive "Open Internet" rules, but rather should promulgate general guidelines to provide notice to industry participants (be they consumers, last-mile or backbone carriers, or edge providers) about the general classes of conduct of concern to the Commission. Should subsequent conduct raise concerns under these guidelines, the Commission should take strong and swift enforcement action, proceeding on a case-by-case basis under Section 706.

This approach, which is captured in the proposed "commercially reasonable" standard, is grounded in modern principles of administrative law and procedure; is sufficient to protect consumers and police potentially problematic behavior; is pragmatic both in that it is less

[1] The author teaches and researches in areas relating to telecommunications law and policy, Internet- and cyber-law, and the regulation of technology; has previously worked in the Federal government in a capacity relevant to the present NPRM; has degrees in both law and economics and has researched and published in technical fields relating to modern telecommunications technology. Further information about the author is at http://law.unl.edu/facstaff/faculty/resident/ghurwitz.

likely to result in further years of Open Internet rulemaking efforts (and concomitant uncertainty) than other alternatives being considered by the Commission and in that it is more likely to survive judicial challenge than those alternatives. It allows for the continued development of new pro-consumer businesses and business models on the Internet while allowing the Commission to investigate and take action against conduct that may harm consumers – in this sense it is unlike other alternatives under consideration that could foreclose the development of pro-consumer businesses and business models in the interest of preventing hypothetical (and often nonsensical) consumer harm. This approach maintains flexibility that is essential to avoid manipulation, capture, and arbitrage by firms that would use the Commission's rules to profit, possibly at the expense of consumers.

These comments proceed in four parts. The first part discusses the paramount importance of consumer welfare in whatever rules the Commission adopts. In addition to common points about the importance of consumer welfare, this discussion argues that consumer welfare urges that the Commission adopt rules that will not result in immediate judicial challenge and additional years of development and litigation, that it requires the FCC to play an affirmative role in educating consumers, and that it requires the development of a dispute and complaint resolution mechanisms that is timely, accessible, and responsive to consumers.

The second part briefly surveys the economics of vertical integration, innovation, and multi-sided markets, to argue that legal standards are preferable to rules. The literature on these topics consistently demonstrate that vertical integration, platform "openness," and multi-sided business models can have inconsistent and unpredictable effects on consumer welfare. Any rule, therefore, that prophylactically mandates or forecloses the adoption of broad classes of business structures or conduct is, almost by definition, arbitrary, and its adoption capricious.

The third part discusses the Commission's authority to implement Open Internet rules. It makes the basic point that adoption of Title II as a legal basis would almost certainly result in additional years of industry and consumer uncertainty as the Commission's rules are challenged by interested parties in court and narrowed through forbearance. This substantial cost comes with little gain, as Title II does not offer substantially more legal authority than Section 706, which is already sufficient to the Commission's legal goals.

The third part's discussion also considers the question of implementation – a focus that has thus far largely been absent from discussion of the Commission's rules. It argues that *ex post* case-by-case enforcement of standards is preferable to the development and enforcement of strong *ex ante* rules, and explains the legal requirements to proceed along such a path.

The final part offers a technical discussion of statistical multiplexing and its relationship to the Notice's consideration of minimum levels of access and nondiscrimination. The level of technical detail is modest, but makes a critical point with respect to concerns about so-called "fast lanes": in statistically multiplexed systems (such as the Internet), when one user (or firm) pays another for committed capacity, all other users of the network benefit when that user's capacity is not being used. Conversely, when one user consumes a statistically disproportionate amount of capacity, all other users on the network are harmed. Understanding the technical principles that undergird how Internet resources are shared is essential to understanding concepts like "best effort," minimum levels of access, and harmful discrimination.

I. First, do no (consumer) harm

It is common – rote bordering on trite – to begin comments such as these by asserting the primacy of the consumer in the regulatory calculus. But it is important to do so. Consumer welfare is the lodestone that should guide all government policy; and, in the form of the "public interest, convenience, and necessity," it is expressly the Commission's guiding principle.

Consumer welfare is important to the Commission's proposed rules in several ways. Foremost, regulation of business models and other conduct online has direct and indirect effects on consumer welfare. Good rules will, on net, increase consumer welfare; bad rules will, on net, reduce it. As is argued later, the economics literature consistently demonstrates that the consumer welfare effects of the sort of business models and conduct addressed by the proposed rules are uncertain – which means that any firm rules adopted today have the potential to harm consumers.

But consumer welfare is also implicated in other ways by the proposed rules – as well as by the broader Open Internet/Network Neutrality debate. Three ways the NPRM affects consumer welfare are discussed here: consumers getting caught in fights between content and distribution; consumer uncertainty and confidence in the industry; consumer access to and understanding of regulatory protections. Each of these concerns should be incorporated into the Commission's analysis and ultimate adoption of any rules.

Consumers often get caught in the middle of fights between content owners and content distributors. We see this today in the fights between Netflix and ISPs, and in retransmission consent disputes. We have also seen it historically, for instance in disputes between long distance and local exchange carriers in the early 1900s, and (more poignantly) between Western Union and the Associated Press in the late 1800s.

Fights between content owners and distributors are often over how to divvy up rents extracted from consumers (i.e., profits). Where this is the case, the actual effect of the dispute on consumers is relatively small – though the resources expended in the fights are deadweight losses, the fight itself is over the distribution of existing rents, not the creation of new rents. This changes, however, when firms involve consumers in these disputes, either by withholding access to consumers or content (e.g., distributor refusing to carry or degrading content or a content provider refusing to provide it), or by involving consumers as parties to the dispute (e.g., by encouraging consumers to petition the government for regulation, based on incomplete or biased information). Importantly, both consumers and government naturally focus their initial scorn on content distributors – perceived as mere middlemen skimming undeserved profit from the valued and desirable wares of the content providers; but it is often later discovered that the content owners are the actual malefactors in the dispute with the distributors.[2]

This demonstrates two important points. First, firms often use the regulatory process for their private gain – and this is often done at the expense of consumers. For this reason, antitrust enforcers are generally reluctant to credit complaints by a firm's competitors, or those who could otherwise gain competitive advantage through regulatory action. And, second,

[2] *See, e.g.,* Gus Hurwitz, *Should we regulate firms we just don't like?*, TechPolicyDaily.com (November 6, 2013) (citing and discussing Menahem Blondheim, *Rehearsal for Media Regulation: Congress Versus the Telegraph-News Monopoly, 1866-1900*, 56 Federal Communications Law Journal 299 (2004)).

consumer effects are often difficult to assess and run contrary to expectations. Again drawing from the antitrust experience, this is the basic reason that most antitrust analysis has moved from *per se* rules, proscribing certain conduct as permissible or impermissible, to Rule of Reason analysis, under which antitrust claims are examined on a fact-specific, case-by-case basis.

These concerns are particularly important in media-related industries, where consumer sentiment may be naturally biased against distributors and in favor of content owners. This suggests in part that the Commission should be cautious in accepting consumer expressions of concern about maintaining the Open Internet, especially where they assert concerns about distributor relationships with content owners as the basis for their concerns.

This also suggests that the Commission incorporate consumer concern into its welfare analysis. The Commission leadership and staff know as well as anyone that the rules at issue in the NPRM have captured substantial consumer attention. Ongoing fights about network neutrality create substantial uncertainty and concern for consumers. No matter whether these concerns, or consumers' suggested approaches to resolving them, are reasonable, they can legitimately be viewed as consumer harm.

There are two things that the Commission should do in response to these concerns. First, in adopting its Open Internet rules, the Commission should consider how the implementation of those rules will affect consumer understanding of and confidence in the industry. For instance, between otherwise equal options, the Commission should favor rules that are more straightforward or timely to implement, that would reduce firms' ability to strand or otherwise strategically use consumers, and that are less likely to immediately end up back in court. This last point is particularly important given the Commission's consideration of using Title II as a basis for its authority. Should it elect to do so, it seems likely that the new rules will quickly be challenged, resulting in several months or years of uncertainty before the rules can be implemented. Even absent such litigation, the Commission will likely face years of discussion and litigation over the forbearance of various Title II requirements. This outcome would constitute substantial consumer harm.

Second, in developing and implementing its rules, the Commission must take seriously the need for consumers to understand the rules, report their concerns, and be provided with explanatory feedback about these concerns. The NPRM considers the creation of an ombudsperson; this proposal should be firmly embraced. A central task of this office should be to educate consumers – particularly by providing substantive explanatory feedback when action is not taken in response to a complaint. As part of this task, information and statistics about complaints and responses to them should be compiled and maintained publicly. It is also important that this office be non-political.

II. Economic and technical complexity urge use of standards over rules

A basic assumption of the NPRM is that it is possible to craft pro-consumer Open Internet rules. While there may be some business practices and other forms of conduct that are sufficiently harmful (or beneficial) to consumers to merit *ex ante* treatment – such as conduct that would be clearly problematic under the antitrust laws – the effects of the wider range of practices or conduct subject to the NPRM is generally ambiguous. As such, the Commission would be ill-advised to adopt strong rules – and such rules would necessarily be arbitrary and capricious. Rather, where the Commission feels that the potential for harmful practices or conduct is great, it should provide guidance as to the standards by which it will assess whether

specific instances are, in fact, harmful. Such guidance is likely a necessary requirement to satisfy Constitutional Due Process and fair notice requirements, particularly where the Commission may want to seek fines or damages against a firm.

Three types of organizational structures demonstrate this ambiguity: "open" innovation platforms, vertical integration, and multi-sided markets.

Net neutrality advocates, the strongest supporters of Open Internet regulation, often assert that the Internet has thrived because it is "open." Indeed, this idea is captured by (or perhaps has captured) the Commission's own caption for this docket: Protecting and *Preserving the Open Internet*. The reality, however, is that "openness" is neither necessary nor sufficient for the sort of growth that the Internet has seen and fostered – indeed, it can limit such growth. Moreover, the Internet never has been open in the way that advocates suggest.

II.a The Internet as an "open" platform

On the first point, literature and experience amply demonstrate that "open" platforms, or general purpose technologies generally, *can* promote growth and increase social welfare; they *can* also create limits.[3] These limits originate along myriad vectors. Developing open systems can impose costs on initial developers, both in their development of the system and their education of third parties in how to interface with it. It can ossify a system, as subsequent changes will have negative spillovers for third parties. It can fragment a system, if subsequent changes are made by the initial developer or subsequent developers choose to fork a project. It can impose technical costs, as an open system may need to be "more robustly" (e.g., inefficiently) engineered in order to support open interfaces. In particular, to be truly "open," a system may need to make internal variables and functions available externally (to those outside the system), even where they are best kept internal to the system[4] – this can lead to substantively inefficient system design, or inefficient use of development resources. An example of this can be seen in the architectural differences between the Internet Protocol stack, which is a four layer model, and the OSI reference model, which is a seven layer model.[5] The OSI model is unquestionable more "open," in that it allows a greatly larger number of interface pairings and the export of a greater number of internal variables and functions to facilitate those pairings; the OSI model, however, is rarely used, largely because it is grossly over-engineered. A protocol designed on the OSI model would cost more to develop, be harder to maintain, and have worse technical performance than one designed on a more streamlined (and less "open") model.[6]

---

[3] The seminal cite is Bresnahan & Trajtenberg , *General Purpose Technologies "Engines of Growth?"*, 65 J. ECONOMETRICS 83, 94–96 (1995). See also Christopher Yoo, *Modularity Theory and Internet Policy*, at 33 (2013) (discussing same), *available at* http://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=1467& context=faculty_scholarship.

[4] See Yoo, at 22–24 (discussing interdependency and information hiding between modules).

[5] The two models' layer approaches represent somewhat different purposes, but are broadly comparable, if on general terms.

[6] As a simple, but important example: the most computationally time-consuming process performed in the TCP/IP stack is the passing the payload (the actual data being sent between endpoints) between layers. In a truly "open" stack, IP requires data to be copied three times across the memory bus; for large packets, this can dramatically increase latency. An integrated, "closed," stack, on the other hand, can perform this same function with a single copy operation. Many high-performance IP stacks do precisely this, often integrating the IP or even TCP processing functions with the network interface ("offloading" the IP or TCP processing). Importantly, this involves a form of deep packet inspection, where the middle layers of the IP stack need to both look down the stack into a packet's

Beyond the academic literature, much of which is grounded on theory or models, the relative merits of open versus closed systems has been vividly played out in practice. The classic fights here have played out between, e.g., IBM (a closed system model) and IBM-compatible PCs (open systems); between Apple (a notoriously closed system) and other OS models (e.g., Linux, an open system; and Microsoft, an intermediary model); between application vendors (e.g., traditional closed-source models and the open-source development model); and between different mobile platforms (e.g., between Apple's closed iOS and Google's partially-open Android platform). The results of these fights are generally well known, and demonstrate the indeterminacy of the value of "openness." Apple is the clear example demonstrating the success of closed models. The Apache web server is a leading example of the value of openness.[7] Google's Android is an enlightening example. Google develops the core OS on an open source basis, but retains closer control over the core suite of applications that run on the platform. One should also consider Wikipedia and its struggles to maintain its quality as an open platform over time.[8]

The conclusion to draw from this discussion is that, while open platforms can generate consumer benefits, there is no *ex ante* reason to believe that on particular platform will do so merely because it is open. To the contrary, closed platforms can prove more valuable than open ones.

The second point to make about the "preservation" of the open Internet is that the Internet simply is not, and never has been, an open platform in the sense that net neutrality advocates take it to be. The TCP/IP stack incorporates many design decisions and compromises that were made that either expressly or incidentally improve or degrade the performance of different types of applications or different users based upon how they connect to the network. Moreover, many of these decisions were made based upon the characteristics of then-available computer hardware – they would likely have been made differently and have had different consequences given today's technology and uses of the network.

This fundamental point has been made clearly and strongly by several research scientists that were intimately involved with the development of the Internet. This is best seen with David Clark. Quoting from a 2009 Communications Daily interview with Clark, Hazlett and Wright recount Clark's description of the early Internet:

> "The network is not neutral and never has been," Clark said, dismissing as "happy little bunny rabbit dreams" the assumptions of net neutrality supporters that there was once a "Garden of Eden" for the Internet. NSFnet, an early part of the Internet backbone, gave priority to interactive traffic, he said: "You've got to discriminate between good blocking and bad blocking."[9]

---

headers and up the stack into system state information, in order to directly place a packet's payload directly into application memory in its initial copy operation.

[7] For important recent work on this topic, see Shane Greenstein and Frank Nagle, *Digital Dark Matter and the Economic Contribution of Apache* (2013) ("We argue that these findings point to a large potential undercounting of the rate or return from IT spillovers from the invention of the Internet, and to a large potential undercounting of 'digital dark matter' in general."), *available at* http://www.nber.org/papers/w19507.

[8] For discussion of the open platforms' questionable reliance on ongoing input and support from third parties, *see* Smith, et al, *Experiences Enhancing Open Source Security in the POSSE Project*, at 4-5 (2003), *available at* http://users.ics.forth.gr/~sotiris/publications/bookchapters/2posse2003.pdf.

[9] Hazlett & Wright, *The Law and Economics of Network Neutrality* (2011) (quoting 2009 Communications Daily discussion with David Clark).

Similarly, Jon Crowcroft explains "the basic realities of the net, which has never been a level playing field for many accidental and some deliberate reasons," concluding that "we never had network neutrality in the past, and I do not believe we should engineer for it in the future either."[10] Others have made this same point from the early years of the FCC's involvement in Network Neutrality.[11]

Much of the discussion about the open Internet – including the proposed rules – focus on openness qua neutral treatment of users, applications, and data by the network. As this is a more constrained understanding of openness than much of the innovation literature cited above, it is useful to note – as Clark, Crowcroft, and others also do – that discrimination in the handling of packets has long been discussed as desirable or necessary in development of the technical standards under which the Internet operates.[12] Similarly, research demonstrates that so-called "non-neutral" treatment can be affirmatively desirable.[13]

II.b. Vertical integration

The literature on vertical integration is related to the literature on innovation and openness – and many of the concerns relating to the NPRM can be expressed in terms of vertical integration. For instance, much of the NPRM is concerned with how a vertically-integrated content & distribution firm treats its own traffic compared to that of its competitors. These are unquestionably valid concerns – indeed, the widest range of potentially problematic

---

[10] Jon Crowcroft, *Net Neutrality: The Technical Side of the Debate*, 1 Int'l J. Comm. (2007).

[11] Douglas A. Hass, *The Never-Was-Neutral Net and Why Informed End Users Can End the Net Neutrality Debates*, 22 Berkeley Tech. L.J. 1565 (2007).

[12] *See, e.g.,* RFC 2475 ("Service differentiation is desired to accommodate heterogeneous application requirements and user expectations, and to permit differentiated pricing of Internet service."); RFC 2638 (discussing paid prioritization, saying: "It is expected that premium traffic would be allocated a small percentage of the total network capacity, but that it would be priced much higher."); RFC 1633 ("real-time applications often do not work well across the Internet because of variable queuing delays and congestion losses. The Internet, as originally conceived, offers only a very simple quality of service (QoS), point-to-point best-effort data delivery. Before real-time applications such as remote video, multimedia conferencing, visualization, and virtual reality can be broadly used, the Internet infrastructure must be modified to support real-time QoS, which provides some control over end-to-end packet delays." ... "The first assumption is that resources (e.g., bandwidth) must be explicitly managed in order to meet application requirements. … An alternative approach, which we reject, is to attempt to support real-time traffic without any explicit changes to the Internet service model. The essence of real-time service is the requirement for some service guarantees, and we argue that guarantees cannot be achieved without reservations. … We conclude that there is an inescapable requirement for routers to be able to reserve resources, in order to provide special QoS for specific user packet streams, or 'flows'.").

[13] *See, e.g.,* Richard T.B. Ma, et al, *On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers*, Procs of 2008 ACM Conf Emerging Network Experiment and Tech (CoNEXT 2008), Madrid, Spain, December, 2008 ("we find the justification of the existence of paid-peering between transit ISPs. ... Our previous work ... showed that ... selfish ISPs have incentives to perform globally optimal routing and interconnecting decisions to reach an equilibrium that maximizes both individual profit and global social welfare. ... In this paper we extend our model ... Our result [finds instances where paid-peering can benefit welfare].") David Clark, *Network Neutrality: Words of Power and 800-Pound Gorillas*, 1 Int'l J. Comm. 701 (2007) ("As a technical mechanism, QoS seems to be beneficial. It directly addresses the real performance requirements of different sorts of Internet traffic … This reality begs the question of whether we can find a set of rules that might distinguish between "good" or "acceptable" forms of discrimination, and "bad" discrimination. Unless we can find a bright line, using regulation of discrimination to define acceptable behavior may cause more trouble than it cures.").

conduct already proscribed by antitrust law, such that further regulation by the Commission would largely be duplicative and generally unwarranted.

Nonetheless, the economic literature of vertical integration provides useful guidance for the Commission to consider. As with the literature on open platforms and innovation, the conclusions relating to vertical integration are consistently ambiguous. Most vertical integration *can* harm consumers; but most vertical integration also *can* benefit consumers. Importantly, in practice firms generally vertically integrate to capture efficiencies, and in so doing such integration ultimately benefits consumers.[14] As explained by FTC Commissioner and Professor of Law and Economics Josh Wright,

> Over a century of antitrust jurisprudence, economic study, and enforcement agency practice have produced a well-understood economic analysis of the competitive effects of a vertically integrated firm's "discrimination" in favor of its own products or services, including widespread recognition that such arrangements generally produce significant benefits for consumers.[15]

One of the most illustrative examples of misplaced concerns about vertical integration leading to harmful regulatory intervention is the Supreme Court's 1948 *Paramount* case, in which the Court broke up the vertically-integrated movie production and distribution industry.[16] Subsequent decades of analysis have not been kind to this action, finding that the vertical disintegration led to, *inter alia*, higher ticket prices for consumers, reduced variety and quality of films being produced, substantial loss of jobs within the movie industry, and, generally, the end of the golden age of Hollywood and the beginning of a multi-decade dark age.[17] One of the central reasons for this is that vertical integration allowed firms to better absorb risk and to better capture the returns on risky investment; conversely, the vertical disintegrated firms faced greater exposure to risk and less ability to capture the rewards of risky investment. The result was a substantial reduction in innovation in the industry: vertical integration can be a key ingredient for basic research, development, and innovation.

II.c. Multi-sided markets

A final related area of literature relates to multi-sided markets: markets in which two or more groups of users interact with one another by means of some platform. The Internet is a prototypical example of a multisided market, with ISPs and backbone providers acting as platforms that facilitate the interaction between end users and edge providers. To not needlessly belabor the point, as with the prior examples, the relevant literature yields consistently inconsistent results. In multisided markets, prohibiting platforms from engaging in discriminatory pricing can yield either consumer benefits or harm.[18]

---

[14] *See, e.g.,* James C. Cooper et al., *Vertical Antitrust Policy as a Problem of Inference*, 23 Int'l J. Indus. Org. 639, 658 (2005) (surveying the literature and finding the vast majority of studies find vertical integration to have procompetitive effects).

[15] Joshua Wright, *Defining and Measuring Search Bias: Some Preliminary Evidence,* Geo. Mason L. & Econ. Res. Paper. No. 12-14, at 5 (2011).

[16] *U.S. v. Paramount Pictures, Inc.*, 334 U.S. 131 (1948).

[17] *See generally* F. Andrew Hanssen, *Vertical Integration during the Hollywood Studio Era*, 53 J. L. and Econ. 519 (2010).

[18] The literature here is voluminous, often demonstrates benefits from non-neutrality, and consistently notes ambiguous results. For some examples (most of which cite to the broader literature) see: Nicholas Economides and Joacim Tåg, *Network neutrality on the Internet: A two-sided market analysis*, 24 Information Economics and Policy 91 (2012) ("We have showed that one can find such parameter ranges both in the monopoly model and in the duopoly model suggesting that network

The examples discussed above make abundantly clear the indeterminacy inherent in any Open Internet rules. It is difficult, if not impossible, to know *ex ante* whether any given business model or type of conduct by online intermediaries will be beneficial or harmful to consumers. Any rules that mandate or foreclose certain practices or conduct on an *ex ante* basis would therefore be, of necessity, arbitrary or capricious. The converse, however, is also true: it is possible that certain practices or conduct can harm consumers, such that the Commission must develop an approach to the regulation of the Internet that allows it to take action where action is, in fact, necessary.

This suggests that the Commission should prefer to adopt general standards over clear rules. The legal basis for such an approach is discussed in the next Part.

Before turning to that discussion, it is useful to urge here that, should the Commission adopt clear rules it must affirmatively embrace and discuss the relevant economic literature to explain why a rule of general applicability is in the public interest. Given the uncertain consumer welfare effects that run throughout the relevant literature, the Commission must affirmatively address these concerns in order to avoid judicial challenge for any rules that it adopts. Indeed, such challenge will likely be warranted, given the arbitrariness such rules would demonstrate. More problematic, such a legal challenge would almost certainly result in further years of consumer doubt and uncertainty relating to network neutrality, which would cause concomitant consumer harm.

III. Legal authority and implementation

A central question raised by the NPRM is whether the FCC should act based upon Section 706 or Title II. This question has been overwrought: by and large, either approach gives the Commission broad and sufficient authority to enforce its proposed rules. There may be some differences on the margin as to what either approach will allow; but by and large, both

---

neutrality regulation could be warranted even when some competition is present in the platform market. *However, the overall effect of implementing network neutrality regulations can still be both positive and negative depending on parameter values*.") (emphasis added); Paul Njoroge, et al, *Investment in Two-Sided Markets and the Net Neutrality Debate*, 12 Review of Network Economics (Feb 2014) ("This paper adds to the growing body of formal economic analysis that will help inform policy makers on the net neutrality debate and sheds light on the validity, or lack thereof, of the arguments proposed by the different advocacy groups involved. In particular, this article develops a game theoretic model based on a two-sided market framework … to investigate the effects of a net neutrality mandate on investment incentives of ISPs, and its concomitant effects on social welfare, consumer and CP surplus, and CP market participation. ... More specifically, the results regarding the comparison between the neutral and non-neutral regimes for our theoretical and numerical-simulation models are as follows. *In both models, the non-neutral regime leads to a higher overall social welfare*. This result is driven by the higher investment levels caused by the non-neutral regime, which in turn increase consumer surplus and CP gross surplus.") (emphasis added); Jay Pil Choi, Byung-Cheol Kim, *Net Neutrality and Investment Incentives*, 41 RAND Journal of Economics (2010) ("Considering all three channels through which net neutrality can have an influence upon short-run total welfare, we can conclude that static welfare implications of net neutrality regulations depend on the trade-off between transportation cost saving and inefficient production. If the margin difference is significantly large relative to the degree of product differentiation, the discriminatory network would be preferred from the viewpoint of social welfare."; "We find that the relationship between the net neutrality regulation and investment incentives is subtle. Even though we cannot draw general unambiguous conclusions, we identified key effects that are expected to play important roles in the assessment of net neutrality regulations.").

approaches offer substantial, but not complete, authority. Given the broad grant of power either of these approaches affords, the more important question is how the Commission's choice of legal basis will impact its implementation of the rules. Implementation considerations strongly urge the Commission both to proceed based upon Section 706, and to implement its rules on a case-by-case enforcement model.

III.a. Section 706 vs. Title II

As upheld by the DC Circuit in January, the Commission has broad authority to regulate the Internet under Section 706.[19] The breadth of this power is largely in the Commission's own discretion, based on principles of deference.[20] The broad scope of this deference is long enduring and has been recently affirmed by the Supreme Court.[21] The primary limit is that the Commission cannot interpret Section 706 to do that which is expressly prohibited by the Communications Act – that is, it cannot regulate information services as common carriers.[22]

Title II is similarly broad. Title II grants the Commission pervasive authority to regulate broad aspects of the telecommunications sector, including myriad statutory sections authorizing specific regulation of prices, interconnection, and business practices – many of which clearly would not or should not fully apply in the Internet context. It is, however, not without limits. In specific, Title II is limited to various "unjust and unreasonable" practices.[23] As made clear previously, the various forms of conduct underlying the concerns the NPRM seeks to address have consistently ambiguous results, such that it is typically unclear, *ex ante*, whether any given conduct is in fact "unjust or unreasonable."

Importantly, clearly expressed statutory language – such as "unjust or unreasonable" – must be given meaning. Congress has very clearly expressed its intent that Title II does not apply to all conduct by telecommunications carriers that falls under Title II; rather, only the narrow class of conduct identifiable as "unjust or unreasonable" can be proscribed by the Commission. In the context of the NPRM, this is an exception that can quickly swallow the whole of Title II authority. Should the Commission choose to go down the Title II path, it will be little more than an invitation for further years of litigation. Curiously, should the Commission choose to adopt strong *ex ante* rules, this suggests that Section 706 may, in practice, give the Commission greater authority than Title II.

The Commission's forbearance authority is of little help in this context. The purpose of Section 10,[24] allowing forbearance of Title II's various requirements, is to allow the Commission to reduce the requirements of Title II in response to changing market conditions – it is not to allow the

---

[19] *Verizon v. FCC*, 740 F.3d 623 (D.C. Cir. 2014)

[20] *Chevron U.S. A. Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837 (1984).

[21] *Arlington v. FCC*, 569 U.S. ___ (2013).

[22] *Verizon* ("Given that the Commission has chosen to classify broadband providers in a manner that exempts them from treatment as common carriers, the Communications Act expressly prohibits the Commission from nonetheless regulating them as such.")

[23] 47 USC 202(a) reads in full:

> It shall be unlawful for any common carrier to make any *unjust or unreasonable discrimination* in charges, practices, classifications, regulations, facilities, or services for or in connection with like communication service, directly or indirectly, by any means or device, or to make or give any *undue or unreasonable preference or advantage* to any particular person, class of persons, or locality, or to subject any particular person, class of persons, or locality to any *undue or unreasonable prejudice or disadvantage*. (emphasis added)

[24] 47 USC 160.

Commission to craft an entirely new regulatory regime. Forbearance was adopted as part of the 1996 Telecom Act's goal of fostering competition in the telecommunications market, recognizing that the justifications for Title II's extensive regulation would diminish over time as entry occurred – and that in some cases Title II could unduly burden potential entrants.[25]

The Supreme Court's recent decision in *Utility Air Group v. EPA* should give the Commission great pause before proceeding down the Title II path.[26] In *Utility Air Group*, the EPA adopted broad regulations over greenhouse gases. In so doing, it acknowledged that the adopted rules were overbroad and administratively infeasible. In order to enforce the new rules as the agency intended, the EPA adopted a "Tailoring Rule" to tailor its enforcement of the regulations to only apply to the conduct the agency intended to target. The Supreme Court rejected this approach, explaining that "agencies must operate within the bounds of reasonable interpretation" and that "an agency interpretation that is inconsistent with the design and structure of the statute as a whole does not merit deference."[27] Like the EPA's regulation of greenhouse gasses, Title II regulation is meant to enable pervasive regulation of clearly problematic conduct, and the clear Congressional intent is that the forbearance mechanism is meant to foster conduct that will render the purposes of Title II nugatory. The Commission could not give meaning to Title II's "unjust and unreasonable" limitations without doing irreparable damage to the underlying statutory structure. The use of forbearance as a "tailoring rule" to structure Title II to suit the Commission's newfound purpose would be "inconsistent with the design and structure of the statute as a whole."

III.b. Implementation of legal authority

No matter whether the Commission acts under Section 706 or Title II, any strong *ex ante* rules that it adopts are likely to be challenged. This follows from the inherent uncertainty as to the effects of any proscribed conduct. Such challenge is particularly likely should the Commission act under Title II, given the concerns described above.

In deciding how to implement its rules the Commission must understand that both the industry and public need resolution over the Commission's Open Internet rules. Ongoing uncertainty over these rules constitutes a real harm to both consumers and industry. And, while telecommunications lawyers enjoy the likely hundreds of millions of dollars that have been spent on net neutrality over the past decade, few likely think this has been money well spent.

The basic problem with adopting strong *ex ante* rules is that whether any specific business model or conduct is harmful is highly fact-specific. Any rules addressing such conduct generally and on prophylactic terms are, therefore, by definition arbitrary and capricious. Importantly, administrative law has developed longstanding principles to address precisely this concern – chief among them agency choice of procedure. It is a longstanding and bedrock principle of administrative law that agencies that have the authority to develop rules of general applicability can develop those rules through case-by-case adjudication. As explained in the Supreme Court's seminal *Chenery II* opinion:

---

[25] *See, e.g.,* Senate Report No. 103-367 (1994) ("S.1822 gives the FCC greater regulatory flexibility by permitting the FCC to forbear from regulating carriers when it is in the public interest. This provision will allow the FCC to reduce the regulatory burdens on new entrants. It will also permit the FCC to reduce the regulatory burdens on the telephone company when competition develops or when the FCC determines that relaxed regulation is in the public interest.")

[26] *Util. Air Regulatory Grp. v. EPA*, ___ U.S. ___ (U.S. June 23, 2014)

[27] *Id.*, slip op. at 10.

> Not every principle essential to the effective administration of a statute can or should be cast immediately into the mold of a general rule. Some principles must await their own development, while others must be adjusted to meet particular, unforeseeable situations. In performing its important functions in these respects, therefore, an administrative agency must be equipped to act either by general rule or by individual order. To insist upon one form of action to the exclusion of the other is to exalt form over necessity.
>
> In other words, problems may arise in a case which the administrative agency could not reasonably foresee, problems which must be solved despite the absence of a relevant general rule. Or the agency may not have had sufficient experience with a particular problem to warrant rigidifying its tentative judgment into a hard and fast rule. Or the problem may be so specialized and varying in nature as to be impossible of capture within the boundaries of a general rule. In those situations, the agency must retain power to deal with the problems on a case-to-case basis if the administrative process is to be effective. There is thus a very definite place for the case-by-case evolution of statutory standards.[28]

The Commission's proposed "commercial reasonableness" standard follows precisely this model – and is both learned and wise. Among the reasons for this are that the Commission clearly does have broad enforcement authority, such that a standards-based approach is less subject to judicial challenge outside of the context of a challenge to actual potentially-harmful conduct, and that any challenge to such conduct will be focused on whether that conduct actually is harmful. If such challenged conduct is actually harmful, the fact of that harm will go a long way toward demonstrating the legitimacy of the Commission's rules. If such challenged conduct is not actually harmful, a judicial loss will largely be cabined, affecting only the specific conduct in question and not jeopardizing the Commission's overall framework.

Should the Commission choose to proceed on a case-by-case basis, it should nonetheless issue general guidelines that provide notice as to the sort of conduct that may or may not raise concerns and the legal basis and metrics by which those concerns will be evaluated. The current proposed commercial reasonableness standard does an admirable job on this front.

Providing such guidelines is necessary to meet Constitutional Due Process and Fair Notice requirements. As explained in the Supreme Court's still recent *Fox II* opinion, "A fundamental principle in our legal system is that laws which regulate persons or entities must give fair notice of conduct that is forbidden or required."[29] Importantly, Constitutional problems arise not "because it may at times be difficult to prove an incriminating fact but rather because it is unclear as to what fact must be proved."[30] In other words, the difficulty of demonstrating that a given business model or conduct is problematic does not mean that the Commission cannot proscribe such conduct, but if the Commission is to proscribe that conduct it must be clear as to how it will evaluate whether the conduct is problematic.

---

[28] *SEC v. Chenery Corp.*, 332 U.S. 194, 202 (1947).
[29] *FCC v. Fox Television Stations, Inc.*, 132 S. Ct. 2307 (2012)
[30] *Id*.

IV. Statistical multiplexing, discrimination, and minimum levels of access

The final subject that these comments address is the relationship between the proposed no-blocking and nondiscrimination rules, and in particular, how prioritization affects "best effort" services and the minimum level of access required for a service not to be blocked. These issues are addressed from a technical perspective, focusing on statistical multiplexing – the basic means by which multiple users and applications are able to share Internet facilities. Importantly, many advocates for strong Open Internet rules assert that prioritization necessarily means a reduction in capacity available for non-prioritized services. From a technical perspective this understanding is simply not correct. This section explains why and offers other policy implications derived from the technical mechanisms controlling how data is sent over the Internet.

IV.a. A brief primer: statistical multiplexing, congestion, and queue management[31]

The Internet is a shared facility. The basic mechanism by which the Internet works is that human-understandable information that is to be communicated between endpoints is digitized and broken into small packets. These packets are sent over a shared communications infrastructure that routes them from a sender to a receiver; the receiver reassembles the packets and converts them back into human-understandable form. This is generally called a "packet switching" network, referring to the use of packetized data to communicate between endpoints.

But the packetization of data isn't the only basic feature of the Internet. Once data is packetized, it needs to be communicated across the network. This is done by a process known as statistical multiplexing. The basic idea behind statistical multiplexing is that a significant portion of any communications channel goes unused by any given application. For instance, a typical voice conversation consists of more silence (gaps between sounds made by the speakers) than speech. Other applications can make opportunistic use of that capacity without adversely affecting the initial application. Statistical multiplexing is the mechanism by which unused capacity by one communications channel is shared by other communications channels. This has two general benefits. First, without statistical multiplexing, five "conversations" would require five communications channels; with statistical multiplexing, those five conversations can be accommodated by two, or perhaps three, communications channels. This reduces cost and complexity of a communications network. And second, this allows individual "conversations" access to much greater capacity during periods of relative quiet on the network.

Packetization and statistical multiplexing are closely related: packet switching is the mechanism that makes statistical multiplexing possible on the Internet. Prior to the advent of packet switching, telecommunications services used either circuit-switched or dedicated lines. Under this model, lines were allocated to specific users for specific periods of time. As a result, they could not be shared between other simultaneous users.[32] Packet switching changed this, allowing for a much finer-grained level of sharing, on a purely statistical basis.

The benefits of statistical multiplexing are immense. Because most communication is "bursty," few applications need the full capacity of a dedicated communications channel.

---

[31] The discussion that follows is meant to give a general sense of complicated technical concepts, to demonstrate their importance to the Open Internet discussion. Some liberties have been taken as to precise technical descriptions in order to make the discussion reasonably accessible to a generalist audience.

[32] This description provides a gloss on the concept of statistical multiplexing. Technically, circuit-switching also is a form of statistical multiplexing, which connections are multiplexed on a much coarser boundary.

Statistical multiplexing therefore allows substantially more aggregate capacity to be offered to a group of users than would otherwise be possible if each user was allocated a dedicated communications channel. Conversely, any additional capacity added to a network to benefit one or a small group of users also creates additional capacity for all other users. Thus, because statistical multiplexing allows users to opportunistically use other users' unused capacity, incremental capacity benefits all users.

Statistical multiplexing is implemented primarily by two mechanisms: the algorithms that individual computers use to determine the rate at which they inject packets into the network ("congestion control" algorithms), and the algorithms that routers use to sort and prioritize packets ("queuing disciplines").[33] Queuing disciplines are particularly important. As routers receive packets from multiple sources they store those packets in temporary buffers before processing them to be sent to their destination. Early routers would process packets in the order that they arrived (i.e., on a first-in-first-out ("FIFO") basis) (fig. 1); if buffers were full, the router would drop any incoming packets (i.e., on a "tail-drop" basis). On naïve inspection, this appears to be a fair (even "neutral"!) approach to handling traffic. In practice, however, FIFO and drop-tail algorithms can result in some streams getting a disproportionate share of available capacity, other streams being locked out of getting any capacity at all, and overall inefficient levels of network utilization.

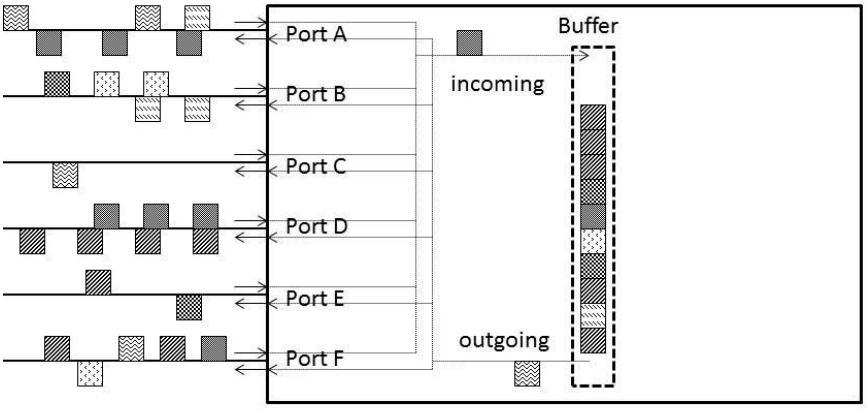

A basic router. As packets arrive they are added to a buffer for processing. The router processes packets one at a time, sending them to a destination port, in the order in which they arrived.

Fig. 1

Given these concerns, engineers have been developing various active queuing disciplines – generally referred to as Active Queue Management ("AQM") – for decades.[34] AQM algorithms allow routers to adopt different approaches to managing queues and dropping packets as buffers reach capacity, in order to ensure "fair"[35] and efficient network operation. This is the mechanism by which "paid prioritization" is most likely to be implemented: ISPs

[33] *See generally* Braden, et. al., *RFC 2309: Internet Performance Recommendations* (April 1998), *available at* http://www.ietf.org/rfc/rfc2309.txt. The term "queuing discipline" used here encompasses both queue management and scheduling algorithms.

[34] *See, e.g.,* Floyd, S., and Jacobson, V., *Random Early Detection gateways for Congestion Avoidance*, 1 IEEE/ACM Transactions on Networking 397 (1993).

[35] "Fair" is a term of art, referring to how congestion control algorithms and router queuing disciplines interact to apportion bandwidth on a shared resource. There are multiple definitions, and no consensus definition. Generally, "fairness" means that, on a connection shared by *N* connections, no connection will prevent any other connection from being able to use at least 1/N of the overall capacity.

can configure their routers to handle prioritized data specially. For instance, prioritized traffic can be moved to the head of the queue. Or, more often, prioritized traffic will be placed in a separate queue, from which packets will be processed more often than from other queues (fig. 2.).[36]

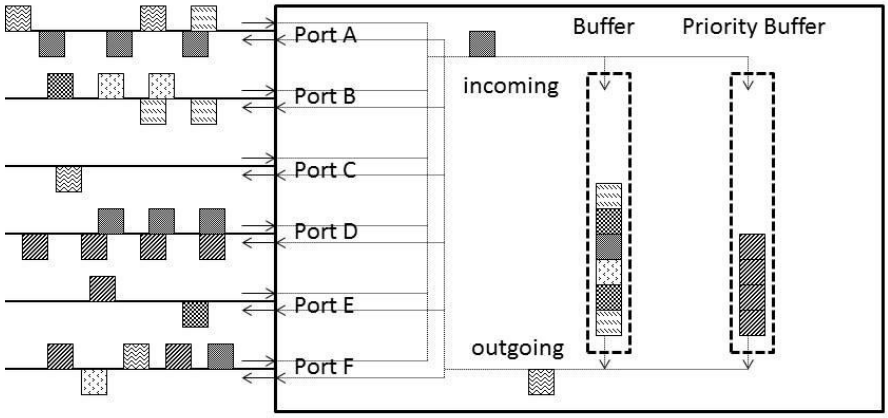

One model of a router with prioritization. As packets arrive, most go into one buffer, while "priority" packets go into their own buffer. The router can process packets from the two buffers in many ways, for instance by alternating between the buffers. This reduces the time that prioritized packets spend "waiting."

Fig. 2

Under any queuing discipline, congestion occurs in one of two ways: packets are either dropped, or they are delayed. Drops occur either when a router's buffer space is full, or when the network connection between devices is at capacity (e.g., trying to transmit data at a rate of 15 mbps for a sustained period over a connection only capable of transmitting 10 mbps).[37]

Delays occur when packets arrive at a router faster than the router can send them to their destination. Where this happens, the router buffers packets until it can process them. Such delays most often occur because Internet traffic is "bursty," meaning that packets from a given sender tend to arrive at a router in groups – not because the destination network has insufficient capacity to handle the data to be sent over it. As an example, if we assume a router with ten 100 mbps ports (and symmetric traffic patterns across all ports), with average

[36] Importantly, queuing disciplines alone cannot efficiently guarantee an equal allocation of bandwidth among streams. Bandwidth allocation is governed by many factors, including congestion on other parts of the network, latency and packet round-trip-time, and congestion algorithm. Certain queuing disciples attempt to ensure equal bandwidth between streams; while such disciplines can increase the share of bandwidth allocated to lower-speed streams (e.g., those that take longer to reach a steady-state packet rate due to relatively higher round-trip times), they often accomplish this at the cost of disproportionately slowing other streams.

[37] In practice, these cases are generally equivalent today. This is because most devices are on one of two types of connection. Either they are on switched ports, meaning that there is only a single sender/receiver pair on each connection. Or they are on an asymmetric shared connection with a single gateway (e.g., the CMTS). In that case, devices on the shared connection may interfere with one another, resulting in packets being sent to the gateway to drop. However, the bulk of data is typically being sent to the devices on the shared connection (e.g., from edge providers, through the CMTS, to end-users). Where this is the case, the gateway is the only transmitted, so drops are likely to result as data from edge providers arrives as a rate exceeding the capacity of the shared connection; in which case, the gateway will buffer packets until its buffers are exhausted, at which point it will begin dropping packets. The great exception to this discussion is wireless networks, where packet drops are more common, for various reasons.

utilization of 40 mbps per port, it is likely that the rate of incoming traffic to each port will regularly exceed several hundred mbps for very short periods of time. Without buffering, this would result in substantial packet loss (i.e., dropped packets); with buffering, the router can queue several hundred or thousand packets (or more) received over, e.g., milliseconds to 10 of milliseconds and send them back out over a period of, e.g., 10s of milliseconds to 100s of milliseconds. Most applications are not particularly sensitive to delays – and even those that are can generally tolerate ther moderate variances that are typically seen on actual networks – so it is often preferable to address congestion with buffering instead of by dropping packets.[38]

Another example is particularly illustrative, as it demonstrates a common instance where buffers may be overrun, resulting in packet loss. The MPEG-DASH protocol, which is used by services like Netflix to stream video with adaptive bitrates, sends video in chunks (commonly in units of 10 seconds of video each).[39] This typically results in an ON/OFF transmit pattern, where video is sent from the server to the client for about 10 seconds, and then sends nothing for about 10 seconds. Researchers are only just starting to understand how this protocol, which transmits for relatively large bursts, interacts with router buffering and TCP's congestion control algorithm. While on average data is being sent to a steady number of clients at a constant data rate, statistically there will be periods where either far more or far

[38] But note, this is not always the case – proper tuning of buffer sizes is a difficult and technical topic. *See, e.g.*, Jim Gettys & Kathleen Nichols, *Bufferbloat: Dark Buffers in the Internet*, 55 Comm's ACM 57 (2012). Additionally, while it may be preferable to address ongoing congestion by buffering instead of dropping packets, congestion *avoidance* is often best implemented by responding to delay. There is longstanding research and debate within the technical community over congestion control algorithms that rely on packet loss (e.g., TCP Reno) vs. packet delay (TCP Vegas). *See, e.g.,* Budzisz, et al., *On the Fair Coexistence of Loss- and Delay-Based TCP*, 19 IEEE/ACM Trans. on Networking 1811 (2011).

[39] For discussions of MPEG-DASH relevant to this and the next paragraph, see: Ahmed Mansy, et al., *SABRE: A client based technique for mitigating the buffer bloat effect of adaptive video flows*, Procs. 4th ACM Multimedia Systems Conf. 214 (2013) ("we show ... that a single DASH stream can cause significant delays to other ongoing applications sharing the home network in a typical residential setti"g."); Jim martin, et at., *Characterizing Netflix bandwidth consumption*, Procs. Consumer Comm's and Networking Conf 230 (2013) ("Ongoing academic research is providing foundations for understanding how DASH applications behave and how they might be improved."; "we seek to understand the impact of adaptive applications on congestion and bandwidth control mechanisms throughout the Internet or within a broadband access network. The work presented in this paper provides foundations for achieving this goal."; "Academic research on DASH is just emerging."; "The dynamics and implications of multiple levels of end-to-end congestion control are not well understood."); Saamer Akhshabi, et al., *An experimental evaluation of rate-adaptation algorithms in adaptive streaming over HTTP*, Procs. 2d ACM Multimedia Systems Conf. 214 (2013) ("Adaptive streaming over HTTP is a new technology. It is not yet clear whether the existing commercial players perform well, especially under dynamic network conditions. Further, the complex interactions between TCP's congestion control and the application's rate-adaptation mechanisms create a 'nested double feedback loop' - the dynamics of such interacting control systems can be notoriously complex and hard to predict."); Te-Yuan Huang, *Confused, timid, and unstable: picking a video streaming rate is hard*, Procs. 2012 ACM Conf on Internet measurement 225 (2012) ("all three [MPEG-DASH] services we study display degraded performance in the presence of competing traffic, well below the video quality possible if the client used its fair share of bandwidth."; "all three services we study display degraded performance in the presence of competing traffic, well below the video quality possible if the client used its fair share of bandwidth."; "In the worst case, the feedback loop creates a "death spiral" and brings the playback rate all the way down to its lowest value."); Junchen Jiang, et al., *Improving Fairness, Efficiency, and Stability in HTTP-based Adaptive Video Streaming with FESTIVE*, Procs. 8th Int'l Conf. Emerging Networking Experiments and Techs. 97 (2012) ("Many commercial video players rely on bitrate adaptation logic to adapt the bitrate in response to changing network conditions. Past measurement studies have identified issues with today's commercial players with respect to three key metrics---efficiency, fairness, and stability---when multiple bitrate-adaptive players share a bottleneck link. Unfortunately, our current understanding of why these effects occur and how they can be mitigated is quite limited.").

fewer than the average number of segments are being simultaneously sent. When more than the average number of segments are being simultaneously sent, given the size of each segment, there is a potential to saturate router buffers and cause packet-loss congestion. Such congestion may give the appearance that the network is under-provisioned by the network operator, when in fact it is largely the result of application-level attempts at congestion control; additionally, while it will adversely affect all traffic being sent over the shared link, it is likely to affect some traffic more adversely than other traffic, potentially giving the appearance of discriminatory treatment.[40]

Critically, from the perspective of the Commission's Open Internet efforts, this demonstrates the limits of our understanding of the basic technologies that the FCC is attempting to regulate. It is difficult to imagine that the Commission can implement an *ex ante* regulatory regime that yields an efficient allocation of these resources given the current state of knowledge about how these resources work. The interaction between MPEG-DASH, router queuing, and congestion control is a current area of cutting edge research. Current research suggests that traditional, best-effort, non-prioritized routing may yield substantially inefficient use of the network resource. It may well turn out to be the case that efficient routing of data like streaming video requires router-based prioritization. It may even turn out that efficient routing of streaming video data is necessarily harmful to other data – it may not be possible to implement a single network architecture that efficiently handles data with differentiated characteristics. If this is the case, then it may certainly be "commercially reasonable" that streaming video providers pay a premium for the efficient handling of their data, in order to compensate for the negative externalities that those uses impose upon other users and uses.

IV.b. Prioritization, Congestion, and Minimum Levels of Access

To date, no firm is known to have offered a paid-prioritization service of the sort contemplated by the NPRM. It is therefore not entirely clear what such a service would look like or how it would be implemented. While necessarily speculative – again demonstrating the danger of adopting strong *ex ante* rules – in all likelihood a paid-prioritization model would be implemented through priority buffering as described above: a firm would pay to have its packets handled on a priority basis by another firm's routers. Taking this as the basic mechanism of prioritization yields several important conclusions about the concepts of minimum levels of access and discrimination – and, importantly, their relationship to one another.

Prioritization can affect connections in one of two ways: by increasing or decreasing that connection's bandwidth, and by increasing or decreasing its delay. The primary effect of prioritization is to decrease a prioritized connection's delay; this follows from the basic prioritization mechanism of handling incoming packets via a priority queue. Unsurprisingly, this can increase the delay to which non-prioritized packets are subjected.

This increased delay[41] is unlikely to meaningfully affect most applications. While some applications are particularly sensitive to delay, delay is a persistent technical reality. All applications are subject to delay, and most are not sensitive to it; all delay-sensitive

---

[40] See, e.g., Te-Yuan Huang, *Confused, timid, and unstable: picking a video streaming rate is hard*, supra (describing the performance "death spiral" that service using MPEG-DASH (such as Netflix) can experience under congestion circumstances).

[41] This discussion does not differentiate between "delay" and "changes in delay" (technically known as jitter). For some applications, especially streaming media, jitter is more problematic than delay. Prioritization would likely yield similar effects for both delay and jitter, though the precise effects are ambiguous.

applications are designed to accept and mitigate delay as best they can. Any increase in delay caused by prioritization is likely to be marginal and fall within the accepted margins of typical delay-sensitive applications. Moreover, because delay is a persistent technical problem that is only exacerbated, not created, by prioritization, a better approach to delay mitigation is the wide-scale implementation of application-controlled or automated prioritization (e.g., QoS) mechanisms by routers. These mechanisms – many of which already exist, but are infrequently implemented – allow routers to prioritize delay-sensitive traffic. Wide-scale implementation of these mechanisms would broadly address any concerns raised by the adverse effects of paid-prioritization on delay-sensitive applications.[42]

Prioritization would also affect how much bandwidth competing streams are able to use. Assuming the network is uncongested – as recent research makes clear is the case on most consumer-Internet backbones – prioritization would allow a prioritized connection to consume more bandwidth. But, counterintuitively and contrary to the understanding of many net neutrality advocates, prioritization would not appreciably affect the bandwidth available to other connections.

This follows from the interaction of queuing disciplines and congestion control in statistically multiplexed systems. Almost every implementation of TCP in use today uses packet loss, not delay, to control its data rate.[43] End-hosts on the Internet do not know how much bandwidth is available to them, so they have no way of knowing at any given time what speed they should send data at.[44] What they do instead is start sending data at a slow speed and incrementally increase that speed until congestion occurs (e.g., packets are dropped). At that point, they know they have exceeded available capacity, so they slow down (usually by reducing transmit rate by half); they then resume incrementally increasing their speed. This process, known as "congestion control" repeats indefinitely.

Because congestion control algorithms primarily rely on packet-loss to signal congestion, so long as a network is not experiencing packet loss prioritization will not substantially affect the rate at which end-hosts send packets.[45] Thus, and directly contrary to

---

[42] In considering the meaning of "minimum levels of access," the Commission may consider, in particular, whether Internet Service Providers implementing paid prioritization should also be required to implement user-selectable or automated QoS, perhaps with a requirement that packets subject to paid-prioritization cannot be given greater priority than applications reasonably recognized as especially delay sensitive. The merits of such a proposal require substantial further investigation, and these comments do not endorse such an approach without further research.

[43] Standard implementations of TCP used by all major operating systems rely on the packet-loss mechanism. Delay-oriented congestion control algorithms also exist, but are less frequently used – largely because they are systematically unable to make as efficient use of the network resource in a system shared with packet-loss oriented TCPs. *See, e.g.,* Budzisz, et al., *On the Fair Coexistence of Loss- and Delay-Based TCP*, 19 IEEE/ACM Trans. on Networking 1811 (2011). The great exception to this is recent versions of Windows, which uses a hybrid loss- & delay-based algorithm ("Compound TCP"). Even under congestion, however, Compound TCP is bounded on the low end by the performance of the standard Reno loss-based algorithm.

[44] This is a fundamental constraint of the Internet. Since no one can know, *ex ante*, how many users will be sharing a connection at a given time, one can never know how much bandwidth is available to a given user at a given time – even if all of the technical characteristics of the network are known (which they are not). For instance, if three users share a 100 mbps connection but only two are active at a given time, a "fair" TCP will let each of them use 50 mbps. If the third user then starts using the connection, only 33 mbps will be allocated to each. There is no way to know whether or when other users will be making use of the shared resource.

[45] The rate at which TCPs increase the rate at which they send data is affected by delay, so there is some marginal effect on speed that results from prioritization. This is, however, a second order effect, so is bounded by a relatively small margin, even in the case of a linear increment. Modern TCPs increasingly use

claims of many advocates, prioritization of some connections does not reduce the capacity available to other connections. Or, if we use the rhetoric of the debate, on connections without substantial packet loss the hypothetical addition of "fast lanes" (prioritized connections) does not relegate other connections to a "slow lane."

IV.c. Implications

This discussion suggests several conclusions relating to the NPRM's consideration of minimum levels of access and prioritization, and in particular how the two relate.

First and foremost, the effects of prioritization should be measured primarily in terms of on-net packet loss, and, in particular, the relative levels of packet loss between prioritized and non-prioritized traffic. If these rates are similar (including the case where they are at or near zero), this strongly suggests that any prioritization is commercially reasonable. Indeed, it strongly suggests that prioritization is not affecting non-prioritized traffic at all. In cases where there is substantial packet loss but the rates are comparable between prioritized and non-prioritized traffic, this suggests the relevant network is under-provisioned – but it does not suggest concerns relating to prioritization or Open Internet principles generally.

The only situation where prioritization on uncongested networks may reasonably be seen as adversely affecting non-prioritized connections is where those connections are substantially delay-sensitive. In such cases, network providers implementing paid prioritization should take steps to ensure such applications are not unduly affected. Such steps could take many forms, from user-selectable prioritization, to automated (e.g., DPI-based) prioritization of delay-sensitive traffic, to other router-based QoS mechanisms (e.g., tuning the rate at which prioritized buffers are serviced relative to non-prioritized buffers). These and similar ideas have long been discussed within the technical community as important to efficient network use; paid prioritization may reasonably increase the urgency of their deployment. Given the (small, but hard to define) range of possible applications that could be adversely affected by prioritization, the uncertain ways and extent to which they may actually be affected in any particular case, and the myriad approaches to mitigating such effects, the Commission should address any concerns arising from prioritization on a case-by-case basis.

In all other cases, the fact that a network is uncongested strongly suggests both that any prioritization is commercially reasonable, and that such prioritization is not adversely affecting other connections (e.g., any minimum level of access that would be provided without prioritization is being maintained on the network with prioritization).

Because a great portion of a given connection's performance is determined by circumstances outside the control of a given network operator (i.e., by off-net factors), competitive benchmarking of services' on-net performance is an important tool in understanding how, or whether, any on-net prioritization is affecting a given service's performance. For instance, as discussed above, the performance characteristics of MPEG-DASH, the protocol currently used by firms like Netflix to deliver streaming video, are not yet well understood. Initial research suggests that MPEG-DASH in particular has poor performance characteristics that are attributable to its design and interaction with TCP's congestion control mechanisms – and that are outside of the control of a given network operator.

---

non-linear recovery mechanisms (e.g., CUBIC, New Reno, Compound TCP), for which a marginal increase in delay is unlikely to substantially affect average transmission rates.

Operationalizing this factor, in evaluating whether prioritization unreasonably affects a given service the Commission should look to the performance of other similar services. If other similar services are able to operate well on a given network, this strongly suggests that any performance issues are related to off-net factors. Even in the absence of such positive comparisons, the Commission should be careful to look at the underlying technical mechanisms and state of relevant research in attributing any performance difficulties to on-net vs. off-net factors.

Related to this point, the earlier discussion of statistical multiplexing leads to a final observation: services that consume a statistically disproportionate amount of the shared network resource (e.g., available capacity) impose a negative externality upon all other users and uses of that resource. Independent from any fees for prioritization, such users should reasonably be expected to pay a premium for their disproportionate use of the resource. A basic premise of the Internet is that aggregation of underutilized resources creates a more efficient aggregate resource. This aggregation is accomplished by statistical multiplexing. But when a single service uses a disproportionate share of the network resource, the opportunistic sharing enabled by statistical multiplexing (i.e., allowing one service to use another service's excess capacity) becomes free-riding instead. Such conduct is more harmful to the Internet's future as an innovative platform than prioritization or other concerns raised in the NPRM; to whatever extent the Commission takes action to police conduct such as prioritization, similar attention should be paid to services that make uncompensated disproportionate use of the network resource.

Conclusion

These comments are respectfully submitted, with the purpose of making clear the need for the Commission to proceed on this matter (to the extent that it proceeds at all) on a case-by-case basis, using administrative adjudication and enforcement actions to address specific conduct in a fact-specific setting. Such an approach is necessary because at an economic and technical level the concerns that have given rise to the Open Internet proceedings have the potential to either harm or benefit consumers. The effects of any particular form of conduct generally cannot be evaluated on an *ex ante* basis, but rather must be evaluated in the context of an actual case or controversy. As such, any strong *ex ante* rule would necessarily be arbitrary and capricious, and otherwise in contravention of modern principles of administrative law and in excess of the Commission's legal authority. Rather, case-by-case adjudication is the typical and preferred approach for Administrative agencies to use in developing legal norms in settings such as this proceeding. Finally, any action that the Commission does undertake, either on a rulemaking or adjudicatory basis, must be based on a sound understanding of the underlying technical principles on which the Internet operates, and with substantial deference to the current state of technical knowledge. Much of the discussion relating to this proceeding, including the basis for many of the underlying concerns, either misunderstands or disregards the underlying technology, often drawing inapt (or inept) conclusions or raising technically implausible concerns. This includes concerns raised by many individuals and firms who work in technology-related fields, but who lack significant expertise relevant to these proceedings. One of the great achievements of the Internet is that one need not understand the vagaries of congestion control or queuing disciplines to create a cutting edge technology firm. Any action undertaken by the Commission must be based, first and foremost, on a sound understanding of the underlying technology, lest the Commission risk both acting in contravention of its legal mandate and significantly undermining the efficient operation of the Internet.

University of Nebraska Lincoln

Justin (Gus) Hurwitz
Assistant Professor of Law
(402) 472-1255
ghurwitz@unl.edu

Before the
Federal Communications Commission
Washington, D.C. 20554

In the Matters of	)	
	)	
Protecting and Promoting the Open Internet	)	GN Docket No. 14-28
	)	
Framework for Broadband Internet Service	)	GN Docket No. 10-127
	)	

Reply Comments of Justin (Gus) Hurwitz, Assistant Professor of Law, University of Nebraska College of Law[1]

These comments are respectfully submitted as reply to comments previously filed by Etsy, Inc, in response to the Commission's May 15, 2014, Notice of Proposed Rulemaking in GN Docket Number 14-28, relating to the Open Internet proceedings, and related matters. They were published on September 15, 2014, as an article on Tech Policy Daily.[2] I hope they may be of help as you continue to consider these matters.

Facts Let us Argue Directly with Comments from Around the World

Recent comments by Etsy CEO Chad Dickerson – as well as comments filed by his company in the FCC's Open Internet proceeding – demonstrate both some of the best and worst of what pro-Title II advocates have to offer. I'll focus here on what makes his comments some of the best. Simply put, they go beyond expressing mere opinion about the FCC's potential regulation of the Internet and attempt to provide a factual basis for that opinion. I happen to find the factual basis provided unpersuasive. But the effort is laudable and deserves serious

[1] The author teaches and researches in areas relating to telecommunications law and policy, Internet- and cyber-law, and the regulation of technology; has previously worked in the Federal government in a capacity relevant to the present NPRM; has degrees in both law and economics and has researched and published in technical fields relating to modern telecommunications technology. Further information about the author is at http://law.unl.edu/facstaff/faculty/resident/ghurwitz.

[2] http://www.techpolicydaily.com/communications/net-neutrality-arguments-facts-figures/

response, as is offered below. If all commentators would approach this process in a similar manner, discussions about the proposed open Internet regulations would be both healthier and more helpful to the FCC officials who are working tirelessly through this rulemaking process.

Two weeks ago, Wired published an editorial[3] by Mr. Dickerson, in which he argued that the Commission's proposed rules "threaten[] any business that relies on the Internet to reach consumers." Most of the piece is blustery rhetoric of the sort typical of most pro-Title II advocacy. But part way through he offers a factual basis to support this view: that "Research from Google and Microsoft shows that delays of milliseconds result in fewer page views and fewer sales in both the short and long term. This is true not just for high-bandwidth services like video, but for any content delivered over the Internet." His company made a similar assertion in the comments[4] it filed with the Commission in July.

Before looking at the substance of this claim, it is worth noting that Mr. Dickerson's editorial would have been just as effective for his intended audience without this assertion. That he chose to provide these numbers speaks well of his guiding principles. Debates such as this are best had with reference to facts and figures, supported by data, and tested by hypotheses.

Despite this praise, I find fault with his and Etsy's use of this data – fault that bears on my own comments[5] filed with the FCC. In my comments, I discussed the technology underlying how packets traverse the Internet, explaining that any "paid prioritization" mechanism is unlikely to have a substantial effect on the bandwidth (i.e., "speed") available to non-prioritized traffic (assuming that the network is not so congested that packets are being dropped). Rather, the effect of prioritization on non-prioritized traffic would likely be limited to marginal increases in latency and jitter, on the order of a few to a few 10s of milliseconds.

On this front, Mr. Dickerson's and Etsy's comments are laudable, insofar as they focus on delay instead of speed. However, as explained below, the sort of delays cited in their comments are unlikely to implicate Open Internet concerns – indeed, it is debatable whether they are as relevant to Etsy's own business as Mr. Dickerson's use of them suggests.

To start, the data cited are not the best available. Indeed, they appear to be selectively chosen and questionably presented. The basic assertion is that "delays of milliseconds" cost companies like Etsy customers, page views, and ultimately revenue. The data supporting this come from blog posts made in 2006 and 2009, covering statements made by Google[6] (2006) or at the Velocity 2009 conference[7]. While these do support the assertion that delay can adversely affect online business, the characterization of the sensitivity in terms of "delays of milliseconds" is disingenuous at best. The examples cited discuss increasing or decreasing page load time by 100s of milliseconds at minimum and in many of the examples by several seconds – in relative terms these deltas often reflect

[3] http://www.wired.com/2014/09/etsy-ceo-to-businesses-if-net-neutrality-perishes-we-will-too/
[4] https://blog.etsy.com/news/files/2014/07/Etsy-Open-Internet-Comments-7.8.14.pdf
[5] http://apps.fcc.gov/ecfs/document/view?id=7521706505
[6] http://glinden.blogspot.com/2006/11/marissa-mayer-at-web-20.html
[7] http://radar.oreilly.com/2009/07/velocity-making-your-site-fast.html

increases or decreases by factors of 2 to 3 (e.g., improving load times from 7 to 2 seconds). While it is not a lie to label any change in page load time in terms of "milliseconds" (which are, after all, units of time), the choice of units is questionable. Mr. Dickerson would have readers (and the FCC) believe that adding 10 milliseconds delay to a page that takes about 2 seconds (2000 milliseconds) to load would put Etsy out of business. Compare this with the actual data he cites, for example that, for Google, increasing page load times from 500 to 900 milliseconds resulted in a 0.59% reduction in searches per user.

There is a more interesting problem, however, with Mr. Dickerson's and Etsy's decision to focus on page load times: while page load time is undoubtedly an important performance metric, it is far from the only one relevant to this discussion, let alone the most important. It is curious to note that all of the data cited are at least half a decade old – from an era when Netflix's main business was distributing DVDs by mail. In more recent years, user-experience metrics have grown increasingly sophisticated. It is common today, for instance, to focus on page "render time" (how long it takes the page to display in the browser vs. how long it takes for the webpage data from the web server to the client), and in specific "above-the-fold" render time (how long it takes for the visible portion of the page to be displayed).

Importantly, these metrics are mostly affected by factors other than basic network performance. Slow-rendering pages, for instance, more likely result from the webpage's design, implementation of cascading style sheets, and use of scripts and other dynamic content than the time it takes to send packets from server to client. It is helpful to note that the results Mr. Dickerson and Etsy cite all measure changes in page-load times that result primarily from how the pages are designed – not from how packets are delivered. (A great example of this is the research cited from AOL, which is the basis for the very useful http://www.webpagetest.org/, which is an excellent – and fun – resource for those who may be interested in the materials discussed here and by Mr. Dickerson and Etsy.) Indeed, as explained in my comments to the FCC, hypothetical "paid prioritization" mechanisms are likely to add perhaps 10s of milliseconds of marginal delay to "non-prioritized" traffic. Even assuming an expertly designed site that minimizes all non-network related delays, it is unlikely that a "non-neutral" network would introduce delay on a scale sufficient to implicate the results that Mr. Dickerson and Etsy cite – and any delay from prioritization is likely dwarfed by other sources of delay.

This point is redoubled by the increase in typical web page size over the past several years – most metrics show an increase from between 100-200 KB before 2010 to over 1 MB today. This means that the typical web page is sent via a stream of 100s to 1000s of individual packets. The relevant metric to understand how long it takes the network to deliver this number of packets is bandwidth (how many bits the network can deliver per second, as moderated by protocol-level congestion control algorithms), not latency (how long it takes a single bit to traverse the network). Again, as explained in my comments to the FCC, in a statistically-multiplexed packet-switched network, prioritization of some traffic is generally unlikely to have any substantial effect of bandwidth available to non-prioritized traffic.

I don't mean to suggest that Mr. Dickerson is wrong that page load times are important to users' online experience, or to the bottom lines of companies like Etsy. Indeed, it is clear that Etsy has invested substantially in developing a high-quality, fast-loading website. It is, however, questionable to suggest that delays from prioritization are likely to implicate these concerns. Prioritization-related delays would likely be at least an order of magnitude below the levels at which Mr. Dickerson's own data suggest concern. Rather, these data demonstrate that firms invest substantially in website design and optimization, and in firm-side infrastructure. A firm may look at these data and decide to pay a team of engineers to redesign its website, spending, for instance, half a million dollars to reduce page load times by 200 milliseconds; it may also look at how much the redesign would cost and decide that the costs don't justify reducing average page load times from 2200 milliseconds to 2000 milliseconds. Prioritization merely presents firms with the same choice of whether the benefits of an investment in improving performance justify the costs of that investment.

I will conclude on a different note. There is substantial room to debate both the facts that Mr. Dickerson and Etsy present and my critique of them. I happen to think that these facts do not support conclusions for which they are being used – but I readily accept that my own interpretations of them can be criticized as well. The important thing is that we can discuss, debate, and test facts and any conclusions drawn from them. We cannot do this with broad assertions about prioritization "destroying the Internet" or putting companies out of business – and this is a criticism that applies to all sides of the Open Internet debate. Mr. Dickerson's and Etsy's comments deserve to be taken seriously, largely because they *can* be taken seriously. When I first read them, they forced me to pause, to return to the relevant research and compare it to my own thinking. That is how the comment process is supposed to work, to assist the Commission in developing informed policy. To the extent that Mr. Dickerson's and Etsy's comments serve to focus our attention on falsifiable facts and testable hypothesis – instead of assertions based on fear and uncertainty – they are a rare and valuable contribution to the debate.

University of Nebraska Lincoln
COLLEGE OF LAW

Justin (Gus) Hurwitz
Assistant Professor of Law
(402) 472-1255
ghurwitz@unl.edu

February 19, 2015

Submitted by ECFS

Marlene H. Dortch, Secretary
Federal Communications Commission
445 12th Street, SW
Washington, DC 20554

Re: Open Internet, GN Docket No. 14-28
Deployment of Advanced Telecommunications Capability, GN Docket No. 14-126

Ms. Dortch,

Please find attached an article, published today, that considers legal challenges to the Commission's ongoing efforts to regulate broadband Internet access: Justin (Gus) Hurwitz, *Regulating the Most Powerful Network Ever*, 10 FSF Perspectives 9 (2015). I am filing copies in both dockets 14-28 (relating to the Open Internet) and 14-126 (relating to Deployment of Advanced Telecommunications Capability).

As always, please let me know if you have any questions, or I can otherwise be of help to the Commission.

Yours,

Justin (Gus) Hurwitz
Assistant Professor of Law
University of Nebraska College of Law

Cc, via e-mail:
Chairman Tom Wheeler
Commissioner Mignon Clyburn
Commissioner Jessica Rosenworcel
Commissioner Ajit Pai
Commissioner Michael O'Rielly
Jonathan Sallet, General Counsel
Stephanie Weiner, Associate General Counsel

P.O. 830902 / Lincoln, NE 68583-0902
(402) 472-2161 / FAX (402) 472-5185

Perspectives from FSF Scholars
February 19, 2015
Vol. 10, No. 9

Regulating the Most Powerful Network Ever

by

Justin (Gus) Hurwitz *

Introduction

At multiple recent events, FCC Chairman Tom Wheeler has emphatically described the Internet as "the most powerful network in the history of mankind."[1] He did so in the context of defending the FCC's myriad regulatory efforts in the broadband space – from redefining "broadband" in the Section 706 Report, to the pending Open Internet order, and the planned pre-emption of state municipal broadband regulations. Speaking about the Open Internet rules, Chairman Wheeler has explained that he "will modernize Title II, tailoring it for the 21st century,"[2] by "taking the legal construct that was once used for phone companies and paring it back."[3]

[1] Remarks of FCC Chairman Tom Wheeler, Silicon Flatirons Center (Feb. 9, 2015), *available at* http://transition.fcc.gov/Daily_Releases/Daily_Business/2015/db0209/DOC-331943A1.pdf; *see also* of same, *available at* https://www.youtube.com/watch?v=vHsHkKpxVkQ (in which Chairman Wheeler was more emphatic than in his prepared remarks); Brian Fung, *FCC chairman warns: The GOP's net neutrality bill could jeopardize broadband's 'vast future'*, WASHINGTON POST (Jan. 29, 2015), *available at* http://www.washingtonpost.com/blogs/the-switch/wp/2015/01/29/fcc-chairman-warns-that-republican-bill-could-jeopardize-broadbands-vast-future/.

[2] Tom Wheeler, *FCC Chairman Tom Wheeler: This Is How We Will Ensure Net Neutrality*, WIRED (Feb. 4, 2015), *available at* http://www.wired.com/2015/02/fcc-chairman-wheeler-net-neutrality/.

[3] PBS Newshour: FCC proposes treating all Internet traffic equally (PBS Television broadcast, Feb. 4, 2015), *available at* http://www.pbs.org/newshour/bb/fcc-proposes-treating-all-internet-traffic-equally/.

With the Commission apparently set upon its regulatory course, attention is turning to whether these regulations will stand up in court. In this *Perspectives*, I lay out some of the judicial challenges the Commission's efforts to regulate the Internet will face, focusing on challenges to regulations under Section 706 and under Title II. To my mind, the legal pitfalls I discuss are serious ones that may well lead to another judicial reversal of the FCC's efforts to adopt net neutrality regulations. And, separate and apart from the ultimate success or failure of these claims, they make clear the rocky road ahead for the Commission – and the uncertainty Chairman Wheeler's path imposes as the Commission "embarks on this multiyear voyage of discovery."[4]

The central question in any challenge is going to be whether the FCC has legal authority to regulate broadband Internet access services – be it in whole or part, under Section 706 or Title II. The FCC is likely to point to *Brand-X*,[5] arguing that the Supreme Court made clear there that the Commission has broad discretion under *Chevron* in how it classifies broadband; and to *Fox I*, in which the Court affirmed agencies' broad discretion to change prior interpretations of the statutes they administer.[6] But while these cases do afford substantial discretion, it is not without limit. Indeed, the Commission's "triple bank shot" theory for justifying regulation under section 706 is just the sort of "Möbius-strip reasoning" that Justice Scalia's *Brand-X* dissent warns "mocks the principle that the statute constrains the agency in any meaningful way." And, more recently Justice Scalia, writing for the Court in *Utility Air Regulatory Group*, reminds us that "agencies are not free to adopt unreasonable interpretations of statutory provisions and then edit other statutory provisions to mitigate the unreasonableness."[7] Yet this is just what Chairman Wheeler plans to do – not just reclassifying broadband, but "modernizing" and "tailoring" Title II in the process.

It is hard to square the application of Title II – even if edited and modernized by the Chairman's forbearing pen – to the Internet, or even to just broadband Internet access services. This is particularly hard to justify following the 1996 Act, which was enacted "to promote competition and reduce regulation,"[8] and which asserts that "It is the policy of the United States … to preserve the vibrant and competitive free market that presently exists for the Internet and other interactive computer services, unfettered by Federal or State regulation."[9] But rather than following the 1996 Act's deregulatory path, the Chairman has assured us that under his plan "there will be ongoing rules in perpetuity."[10]

Litigation over the Chairman's proposal will focus on a number of specific arguments, such as those considered below. But the FCC should lose on this general argument alone: the Internet indeed is, as Chairman Wheeler says, the most powerful network in the history of humankind –

---

[4] *Util. Air Regulatory Grp. v. EPA*, 134 S. Ct. 2427, 2446 (2014) (*UARG*).
[5] National cable & Telecommunications Assn. v. Brand-X Internet Services, 545 U.S. 967 (2005) (*Brand-X*).
[6] *FCC v. Fox Television Stations, Inc.*, 556 U.S. 502, 515 (2009) (*Fox I*)
[7] *UARG* at 2446. (quoting Judge Kavanaugh, "Allowing agencies to exercise that kind of statutory re-writing authority could significantly enhance the Executive Branch's power at the expense of Congress's and thereby alter the relative balance of powers in the administrative process. I would not go down that road.")
[8] Pub. L. 104-104 (Feb. 8, 1996)
[9] 47 USC 230.
[10] PBS Newshour, *supra* note 3.

and we expect Congress to speak clearly if it wishes to assign an agency to regulate things of such vast economic and political significance.[11] The contortions of "triple bank shots" and "modernizing" Title II and "tailoring it to the 21st century" demonstrate the illegitimacy of the Chairman's chosen path. Even if forbearance allows the Commission to step back parts of Title II, the need to "rewrite clear provisions of the statute should have alerted [the FCC] that it had taken a wrong interpretive turn."[12] This is just another example of Justice Scalia's recent lament: "Too many important decisions of the Federal Government [that] are made nowadays by unelected agency officials exercising broad lawmaking authority, rather than by the people's representatives in Congress."[13]

The rest of this *Perspectives* considers in more specific detail the Commission's claims of authority under Section 706 and Title II.

A few notes at the outset. First, I will necessarily address these challenges holistically. Key portions of the FCC's regulations have yet to be released; more important, all these efforts are intertwined. Second, a number of likely challenges are not considered here, including procedural challenges (e.g., whether the proposed Open Internet rules are a logical outgrowth of, or otherwise were sufficiently noticed by, the May 2014 NPRM); preemption challenges (i.e., whether the FCC can preempt state broadband regulations under *Nixon*); the legal basis for the Commission's regulation of CDNs and interconnection agreements, and generally of those to whom the services are not being provided "for a fee"; the classification of mobile data services under Title II; as well as constitutional arguments (takings and First Amendment) – among other possible challenges. The Commission's many attempts to regulate broadband will face a truly staggering number of legal challenges. And third, speculating on legal outcomes is a risky business. I expect that the FCC faces likely, and substantial, losses as these regulations move to and through the courts – but it is entirely possible that the FCC could prevail on any or all of these claims. Of course, this litigation uncertainty nevertheless gives lie to Chairman's Wheeler's assertion that certainty is the great virtue of the proposed rules.[14]

The Section 706 Arguments

The natural starting place is Section 706. Since the FCC's initial classification of cable modem services, and subsequent classification of DSL, under Title I, Section 706 has been the front line in the FCC's efforts to regulate Internet access. The FCC's 2010 re-interpretation of Section 706 as an independent grant of authority is a monumentally important jurisdictional claim that runs through all of the FCC's current efforts. And it is already before the courts, with US Cellular's *cert* petition in its challenge to the USF/ICC order currently pending before the Supreme Court.

[11] Paraphrasing Justice Scalia in *UARG*.
[12] *UARG* at 2446.
[13] E.P.A. v. EME Homer City Generation, L.P., 134 S.Ct 1584, 1610 (2014) (Justice Scalia dissenting).
[14] *See, e.g.*, C-SPAN Communicators: Communicators with Gigi Sohn (C-SPAN, Feb. 6, 2015), *available at* http://www.c-span.org/video/?324180-1/communicators-gigi-sohn ("Peter Slen: What's the advantage to a Comcast, an AT&T, a Verizon in your view, if this proposal goes through? Gigi Sohn: Certainty.")

Jurisdiction under Section 706

The first question is whether Section 706 grants the Commission affirmative authority to adopt regulations pertaining to the Internet. This is primarily a question of *Chevron* deference: the language of Section706 does not speak directly to the Internet, but the FCC has adopted a construction of the section that brings the Internet under its ambit. The question is whether this is a permissible reading of the statute. In January and May of last year, the DC Circuit and 10th Circuit, respectively, found that this is a permissible construction.

At the time of these cases both courts had sound basis for their decisions. The Supreme Court's recent *Chevron* jurisprudence has been as deferential as ever to agency constructions of their ambiguous statutes. Indeed, the FCC has been at the forefront of the cases establishing the low bar for deference. For instance, in *Fox I*, the Court held that agencies are largely free to change their prior constructions of an ambiguous statute, even where the new construction directly contradicts the prior one.[15] And in *City of Arlington* the Court held that agencies receive the same deference over jurisdictional questions as over substantive questions.[16]

Neither of these opinions was unanimous, however, with different coalitions of Justices dissenting in each.[17] And, while they provide a basis for the DC and 10th Circuit's decisions relating to Section 706, the Court's latest relevant case – *Utility Air Regulatory Group* (*UARG*), decided just one month after the 10th Circuit's ICC/USF decision – suggests the Commission is facing a gusty wind of change.

In *UARG*, the Court rejected an otherwise permissible construction of one statutory provision that was only feasible when taken in conjunction with a separate, problematic, construction of another statutory provision. This case will be discussed in more detail below. For the purposes of Section 706, what is important is the Court's strong reaffirmation of its holding in *Brown & Williamson.*[18] The Court reminds us that it "expect[s] Congress to speak clearly if it wishes to assign to an agency decisions of vast 'economic and political significance.'"[19] It also reminds us that it is a "fundamental canon of statutory construction that the words of a statute must be read in their context and with a view to their place in the overall statutory scheme."[20]

This language speaks clearly as a limitation on the use of Section 706 as an affirmative grant of authority to regulate the Internet – more clearly than the use of general principles of *Chevron* deference to support such a use. A decision to regulate the Internet is, without question, one of "vast economic and political significance." In the Chairman's own words, after all, it is "*The*. Most. Powerful. Network in the history of mankind." It is "America's most important platform for economic growth, innovation, competition, [and] free expression … [and] has been, and remains to date, the preeminent 21st century engine for innovation and the economic and social

[15] *Fox I*.
[16] *City of Arlington v. FCC*, 133 S.Ct. 1863 (2013).
[17] Justice Scalia wrote for the majority in both opinions. Justices Stevens, Ginsburg, Breyer, and Souter dissented in *Fox I*; Chief Justice Roberts and Justices Kennedy and Alito dissented in *City of Arlington*.
[18] *FDA v. Brown & Williamson Tobacco Corp*., 529 U.S. 120 (2000).
[19] *Id*. at 133.
[20] *UARG* at 2441.

benefits that follow."[21] Congress has never given the FCC, or any agency, such power over something of such "vast economic and political significance" – and, to the extent that one can claim that the Commission's authority over telephone networks bootstraps such authority, such a basis is not Congress speaking clearly. To the contrary, Section 706 is more naturally read as being deregulatory; it is part of a statute with an affirmative deregulatory goal; and, to the extent the statute does state a policy with respect to the Internet, that policy is to preserve an Internet "unfettered by Federal or State regulation."[22]

The FCC's "triple bank shot" theory gives lie to the absurdity of its jurisdictional claim. While perhaps a permissible construction in a world of loose deference, it is the very sort of "Möbius-strip reasoning" that Justice Scalia warns in *Brand-X* "mocks the principle that the statute constrains the agency in any meaningful way." The statute requires deregulation, and the agency responds that deregulating requires more regulation. Up, suddenly, is down; the floor is on the ceiling; and words have no meaning. Again, as Justice Scalia says, "This is a wonderful illustration of how an experienced agency can (with some assistance from credulous courts) turn statutory constraints into bureaucratic discretions," and proves Justice Scalia's concern that "Too many important decisions of the Federal Government are made nowadays by unelected agency officials exercising broad lawmaking authority, rather than by the people's representatives in Congress."

Substance of Rules under Section 706

Although most attention has been paid to the FCC's use of Section 706 as an affirmative grant of regulatory authority, any application of that authority is also subject to challenge. Here, too, the Commission faces likely trouble.

The Commission's Open Internet rules, for instance, will surely be challenged as "arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law" (more simply referred to as "arbitrary and capricious"). The basic test the courts use is laid out in *State Farm*, which directs the courts to consider whether, in making its decision, the agency relied on improper factors, failed to consider factors required by its statute, adopted rules contrary to the supporting facts, or otherwise adopted a decision so implausible that it cannot be ascribed to mere difference in views or agency expertise.[23]

While we don't yet know the details of the FCC's Open Internet rules, there is substantial reason to believe that they will fail this test. We can see this, for instance, in the announced ban on paid prioritization. There is substantial economic and technical literature that makes clear both that paid prioritization can be beneficial to consumers and otherwise support to goals of Section 706, and, conversely, b) that banning paid prioritization can be harmful to consumers and detrimental to the goals of Section 706. It is hard to see how an absolute ban on paid prioritization can be anything other than arbitrary and capricious. Indeed, discovery could well demonstrate the extent

[21] NPRM, *In the matter of Protecting and Promoting the Open Internet* (May 15, 2015), *available at* https://apps.fcc.gov/edocs_public/attachmatch/FCC-14-61A1.pdf.
[22] 47 USC 230.
[23] *Motor Vehicle Mfrs. Ass'n of the U.S., Inc. v. State Farm Mut. Auto. Ins. Co*., 463 U.S. 29, 43 (1983).

to which such a decision is based on political concerns – the ideological beliefs of FCC officials or support of the President's efforts to establish his legacy[24] – concerns that are clearly "improper factors."

Application of Section 706

We can see more clearly the challenges that the Commission is likely to face by looking at the recent Section 706 Report.[25] This report, in which the Commission changed its definition of "broadband Internet" from 4Mbps/1Mbps (down/up) to 25Mbps/3Mbps, will likely be the basis for various regulatory efforts that the Commission plans to undertake.[26] Such efforts already include, for instance, the Commission's ICC/USF order (which is based on the extension of Section 706 as a grant of regulatory authority) and planned preemption of state regulation of municipal broadband (the need for which will be based on allegedly insufficient deployment of broadband Internet access under the 706 Report's definition).

The 706 Report has a number of glaring problems, which will poison any decisions made based upon it. For instance, in adopting its 25Mbps/3Mbps broadband definition, the Commission considered, and chose between, only two data points: 10 Mbps and 25 Mbps (both downstream). As explained by the Commission, "We have data for 10 Mbps downstream and 25 Mbps downstream but nothing between those speeds."[27] This recalls the joke about the economist searching for his keys under the streetlight – but, sadly, this is not a joke. A lack of available data does not justify making what is an exceptionally aggressive determination based upon sparsely available facts.

More problematic is the 706 Report's failure to consider prices – or anything other than speed. Consider, the Report explains that it "*must* look at a variety of factors that affect access to broadband" (emphasis added), "includ[ing] an assessment of a variety of factors indicative of broadband availability, such as price"[28] – and, in fact, the Report says that "For purposes of evaluating broadband availability, we examine not only physical deployment and adoption, as presented above, but also quality and price."[29] Yet, despite these assertions that the Report both *must* and *does* consider price, the Commission is "unable to provide any rigorous analysis regarding price"[30] so instead "measure[s] advanced telecommunications capability in terms of speed only."[31] This is, by definition, arbitrary and capricious: the FCC finds that it *must* consider a factor, but then makes a decision that expressly does *not* consider that factor.

[24] Gautham Nagesh & Brody Mullins, *Net Neutrality: How White House Thwarted FCC Chief*, WALL ST. JOURNAL (Feb. 4, 2015), *available at* http://www.wsj.com/articles/how-white-house-thwarted-fcc-chief-on-internet-rules-1423097522.
[25] 2015 Broadband Progress Report and Notice of Inquiry on Immediate Action to Accelerated Deployment, GN Dkt. No 14-126 (Jan 29, 2015) (hereafter *706 Report*)
[26] *See id*, paras. 153–163.
[27] *Id*. para 48.
[28] *Id*. para 65.
[29] *Id*. para 100.
[30] *Id*. para 100.
[31] *Id*. para 25.

But there is an even deeper problem with the Commission's 706 Report: its focus on "high-quality" services at the exclusion of considering the "availability" and "deployment" of those services. The problem can be stated intuitively: because the Commission considers only two different speed tiers and does not consider prices, it implicitly expresses a preference for a world in which, for example, 40% of rural Americans have access to 25 Mbps Internet service for $80/mo over a world in which 80% of rural Americans actually purchase 10 Mbps Internet service at $40/mo. Most, I would posit, would find the latter a better world – and, for the purposes of the statute, one in which high-quality broadband services were available and deployed to more Americans. But whether that is the case or not, the fundamental problem is that the FCC does not even consider it as an option. This is particularly troubling given that the statute is clear that its focus is on the deployment of high-quality service to "all Americans," not on the deployment of highest-quality services to those willing or able to pay for it.

This, despite the FCC's own data suggesting the importance of and variation in prices. The 706 Report notes multiple times that price is an important factor cited by consumers in their decisions whether to adopt broadband. Making 25 Mbps broadband available at a high price does nothing to promote uptake of that service; it may well be far better to promote the availability of lower-speed, much lower-cost, broadband. Indeed, the International Bureau's International Broadband Data Report (IBDR),[32] adopted on the same day as the 706 Report, contains important relevant data – which the Commission did not consider in its 706 Report. The IBDR data shows, for instance, that the United States has among the lowest priced broadband availability when measured in terms of cost per gigabyte.[33] And compared to the EU27+4, the US has significantly better high-speed rural coverage: 44.6% vs. 13.2% in 2012. These numbers are likely related: consumers' price sensitivity suggests that promoting low-cost, moderate-speed broadband promotes adoption that, in turn promotes deployment of higher-speed services. Troublingly, the FCC did not consider any of these factors – it did, however, cherry-pick data from the IBDR that buttress its 706 Report conclusions.[34]

The Title II Arguments

While past and currently pending challenges have focused on the Commission's authority under Section 706, the Chairman's expected reclassification of broadband Internet access services as telecommunications under Title II will raise a slew of additional challenges. The first question is whether the Commission, having previously determined that these services do not fall under Title II, can *change* that classification. The second question is whether the FCC has the authority to subject these services to Title II regulation. And the third question considers the substance of the specific regulations that the Chairman's plan would impose.

Reclassification

The most familiar question relating to Title II is whether the FCC, having previously classified Internet access as a Title I "information service" can now reclassify it as a Title II "telecommunications service." Proponents of Chairman Wheeler's plan have made clear their

[32] Fourth International Broadband Data Report, GN Dkt. No 14-126 (Jan 29, 2015).
[33] *See id.*, tables 4o, 4p.
[34] *706 Report*, paras. 8, 130-132.

confidence that the courts will uphold such reclassification. They point to *Fox I*, in which the Court held that an agency's prior construction of its statute generally does not create obstacles to a changed construction of that statute so long the changed construction is otherwise permissible. And they point to *Brand-X*, in which all nine justices seemed amenable to classifying Internet access as a telecommunication service.

This is not an unreasonable understanding of the law – but it is also not complete, and the outcome is not as certain as its proponents suggest. While agencies' discretion to change between otherwise-permissible constructions of the law is broad, it is not unbounded. As explained in *Fox I*, for instance, reliance interests may create a heightened bar to a changed interpretation.[35] The meaning of this limitation has not been developed by the courts, but there is a sound argument that an industry that has invested hundreds of billions of dollars in private capital based on the prior classification has an established reliance interest. Similarly, changing the classification raises serious constitutional due process and takings concerns.

Fox I also explains that an agency changing its evaluation of facts may justify greater scrutiny of the new policy.[36] In the case of reclassification, the Commission will need to explain why its prior determination regarding incorporation of features such as DNS service does not preclude classification as a telecommunications service. The may be difficult given that ISPs today have become far *more* involved in managing the traffic flowing to and from users' computers – including, for instance, filtering harmful traffic. As a result of these (overwhelmingly pro-consumer) services, users have less ability to control what information is sent from or received by points of their specification – the basic requirement for a service to be a "telecommunications service."

The Commission will also have to address the criticism that its existing policy has been overwhelmingly successful, and the changed policy is not in response to any manifest harms. The Internet economy has thrived under the Title I model, especially in the United States. Indeed, it is well documented that investment in the US has substantially outpaced that in Europe and the rest of the world. In response to capital flight from Europe to the US, European leaders are consistently calling for a more deregulatory approach to the Internet – including a growing skepticism towards the sort of rules that Chairman Wheeler is proposing here. Had there been examples of clear problems under the Commission's Title I approach to the Internet, it would be much easier to meet *Fox I*'s requirement to justify its changed assessment of the facts – instead, given the overwhelming success of its prior approach the Commission is likely to face at least some resistance in these efforts.

Neither of these arguments is dispositive – and, importantly, the contours of the relevant legal standards have not been meaningfully developed through litigation. While it is not unreasonable

[35] *Fox I*, at 515 (the agency may need to "provide a more detailed justification than what would suffice for a new policy created on a blank slate … when its prior policy has engendered serious reliance interests that must be taken into account.").

[36] *Id.* (the agency may need to "provide a more detailed justification than what would suffice for a new policy created on a blank slate … when, for example, its new policy rests upon factual findings that contradict those which underlay its prior policy. … In such cases it is not that further justification is demanded by the mere fact of policy change; but that a reasoned explanation is needed for disregarding facts and circumstances that underlay or were engendered by the prior policy.").

to think that the Court will uphold the Commission's changed classification of broadband Internet access from an "information service" to a "telecommunications service," it also is not as foregone a conclusion as Open Internet advocates assert.

Application of Title II to Broadband Internet Access

Separate from whether the Commission can *change* its prior classification of broadband Internet access services is whether classifying such services under Title II is appropriate at all. Proponents of the Chairman's proposal take for granted that such classification is appropriate. This stems largely from the apparent support that all nine Justices expressed for such a classification in *Brand-X*. That confidence is misplaced. Whether it is appropriate to classify broadband Internet access as a telecommunications service was not a question at issue in *Brand-X*, such that the Court did not inquire deeply into, let alone decide, the matter. Even if it had, the Internet today – both in terms of how it is accessed and its importance as a social, economic, and political platform – is different today than it was a decade ago. And, perhaps most important, the framing of the question, as presented by the Chairman, is different today than it was previously.

As with Section 706 and reclassification, the Commission will receive the benefit of strong deference. But deference is not unlimited. *Chevron*'s basic inquiry is into Congressional intent: did Congress intend for the Commission to exert the authority that it is claiming? Statutory ambiguity, such as exists in the Communications Act's definitions, is an important condition for an agency to claim deference – lacking such ambiguity, the intent of Congress is generally clear, and "that is the end of the matter; for the court, as well as the agency, must give effect to the unambiguously expressed intent of Congress."[37] But ambiguity is only the start of the matter – ambiguous terms alone do not render Congressional intent unclear; nor does it give the agency carte blanche in resolving the ambiguity. As explained by the Court in *UARG*:

> [R]easonable statutory interpretation must account for both "the specific context in which . . . language is used" and "the broader context of the statute as a whole." A statutory "provision that may seem ambiguous in isolation is often clarified by the remainder of the statutory scheme . . . because only one of the permissible meanings produces a substantive effect that is compatible with the rest of the law." Thus, an agency interpretation that is "inconsisten[t] with the design and structure of the statute as a whole" does not merit deference.[38]

UARG frames the basic challenge to classification of broadband Internet access services as telecommunications services. The Chairman's plan involves not just classifying these services under Title II, but "modernizing Title II, tailoring it for the 21st century," by "taking the legal construct used for phone companies and paring it back to modernize it." In effect, the Chairman means to create a new legal regime for the regulation of broadband Internet access. Indeed, it almost necessarily *must* create a new legal regime, because applying some portions of Title II does not make sense in the context of the modern Internet. The Court is likely to view efforts to

[37] *Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837 (1984).
[38] *UARG*, at 2442.

rewrite the statute in this way with substantial skepticism, because the need to "modernize" the statute suggests that the Chairman's proposal is incompatible with the Congressionally-designed statutory structure.

In *UARG*, the EPA's decision to regulate greenhouse emissions for motor vehicles triggered statutory permitting requirements for stationary sources of greenhouse gasses as well. The statute requires permitting of any stationary source emitting more than 250- (or in some cases 100) tons of air pollutants per year. Classifying greenhouse gasses as air pollutants substantially increases the number of stationary sources subject to this permitting requirement – a burden that neither the EPA nor those subject to its regulations could reasonably be expected to meet. In order to avoid this absurd result, the EPA adopted a "tailoring" rule, under which it would only enforce the permitting requirement for stationary sources emitting 100,000 tons of greenhouse gasses per year.

The Court found fault with, and rejected, this approach on several grounds. First, as explained by Justice Scalia, entirely separate from the need to "tailor" the statute the fact that the EPA's adopted approach "would place plainly excessive demands on limited governmental resources is alone a good reason for rejecting it."[39] We do not expect that Congress intended to place implausible burdens on an agency (or, for that matter, those it regulates), especially if alternative readings of the statute are possible. Second, the Court finds that the "EPA's interpretation is also unreasonable because it would bring about an enormous and transformative expansion in EPA's regulatory authority without clear congressional authorization."[40] Both of these concerns reflect the basic premise that the Court "expect[s] Congress to speak clearly if it wishes to assign to an agency decisions of vast 'economic and political significance.'"

There is no question that applying the full force of Title II – a statute designed to regulate an industry that effectively consisted of a single firm – to an industry that today comprises literally thousands of firms[41] would impose excessive burdens on both the FCC (and state regulators) and the industry. Title II, for instance, *requires* the Commission to examine every detail of a telecommunications carrier's business, including all transactions that relate to "the furnishing of equipment, supplies, research, services, finances, credit, or personnel." Title II gives the Commission authority over, and requires in the first instance that it exercise authority over, every aspect of a regulated carrier's business. This may have made sense in the era of a telephone monopoly, especially one using relatively simple technology deployed in a static manner to provide a small number of services to homogeneous customers. In makes absolutely no sense in the context of today's market – a market in which thousands of firms offer service using myriad technologies via networks that are effectively rebuilt every 18-24 months to support a vast array of consumer uses under many different business models. The burden of applying Title II would clearly be excessive and would clearly amount to an "enormous and transformative expansion" of the Commission's authority.

[39] *UARG* at 2444
[40] *Id.*
[41] For instance, WISPA (which opposes Chairman Wheeler's plan) has nearly 800 Wireless ISP members, the American Cable Association (which also opposes Chairman Wheeler's plan) represents over 800 small and medium cable-based ISPs, and there are at least another 700 LECs offering broadband Internet access.

That the Chairman's proposal would be an "enormous and transformative expansion" in the Commission's authority can be seen in other ways, as well. Despite his strenuous efforts to maintain the contrary, the Chairman is proposing to regulate what he has called on several recent occasions "the most powerful network in the history of mankind."[42] And, as the first sentences of last May's NPRM tell us, "The Internet is America's most important platform for economic growth, innovation, competition, [and] free expression … [It] has been, and remains to date, the preeminent 21st century engine for innovation and the economic and social benefits that follow."[43] The telephone network, which Title II was designed to regulate, is an important social and economic tool – but it would be a far stretch indeed to call it "the most important network in the history of mankind," or to call it "America's most important platform for economic growth, innovation, competition, [and] free expression." The Communications Act authorized the Commission to regulate the former, not the latter.

The proposal is also contrary to the regulatory structure and approach that Congress has laid out for the Internet. While the Communications Act, as amended by the 1996 Telecom Act, is almost silent with respect to the Internet, it is not entirely silent. Section 230 asserts very plainly that "It is the policy of the United States … to preserve the vibrant and competitive free market that presently exists for the Internet and other interactive computer services, *unfettered by Federal or State regulation*."[44] This is in line with the overall structure and purpose of the Telecom Act, which was enacted "to promote competition and *reduce* regulation."[45] The FCC's approach to the Internet has followed this deregulatory path – and over the past 20 years Congress has not seen a need to change either the statute or the Commission's path. Indeed, one of the most important things that Congress has done legislatively with respect to the Internet is to consistently prohibit federal or state taxation of Internet access. Yet the Chairman has expressly acknowledged that his proposal would permit collection of Universal Service fees on broadband Internet access. While the mechanism is different from taxation, the effect is the same – an incremental increase in costs imposed by the state – and that effect is contrary to nearly two decades of clear Congressional policy.

This brings us to the EPA's "tailoring rule" and its relation to the Commission's forbearance authority. In *UARG* the Court found that the EPA's efforts to avoid the excessive burdens resulting from its statutory construction by "tailoring" the permitting requirements were problematic for two reasons. First, the tailoring effort was itself problematic; and second, the need to tailor the rules demonstrated that its interpretation was contrary to the statutory scheme designed by Congress. In discussing these concerns, the Court explained:

> An agency has no power to "tailor" legislation to bureaucratic policy goals by rewriting unambiguous statutory terms. … Were we to recognize the authority claimed by EPA in the Tailoring Rule, we would deal a severe blow to the Constitution's separation of powers. … The power of executing the laws necessarily includes both authority and responsibility to resolve some questions

[42] *See supra* note 1.
[43] NPRM, *In the matter of Protecting and Promoting the Open Internet* (May 15, 2015), *available at* https://apps.fcc.gov/edocs_public/attachmatch/FCC-14-61A1.pdf.
[44] 47 USC 230 (emphasis added).
[45] Pub. L. 104-104 (Feb. 8, 1996) (emphasis added).

> left open by Congress that arise during the law's administration. But it does not include a power to revise clear statutory terms that turn out not to work in practice. …
>
> We reaffirm the core administrative-law principle that an agency may not rewrite clear statutory terms to suit its own sense of how the statute should operate. … Instead, the need to rewrite clear provisions of the statute should have alerted EPA that it had taken a wrong interpretive turn. … Because the Tailoring Rule cannot save EPA's interpretation of the triggers, that interpretation was impermissible under *Chevron*.

The Chairman's proposed "modernization" of Title II is clearly an effort to "revise clear statutory terms that turn out not to work in practice" and to "rewrite clear statutory terms to suit its own sense of how the statute should operate." This suggests the Commission has "taken a wrong interpretive turn," and that its "interpretation [is] impermissible under *Chevron*."

The obvious response to these concerns is that, unlike the EPA, the Communications Act gives the FCC express power to forbear from enforcing unnecessary provisions of the Act (including portions of Title II).[46] As such, the Commission is Congressionally-authorized to forbear from applying those portions of Title II that would cause the statute "not to work out in practice." The Chairman's proposed approach would therefore be in line with Congressional design, Congress having expressly granted forbearance authority. This is a legitimate and important difference between the EPA's and FCC's tailoring efforts. But the Commission's forbearance authority is not unlimited, and reliance on forbearance to save the proposed rules is arguably subject to same infirmities as the EPA's approach.

There are at least two reasons why forbearance does not save the FCC's tailoring approach. First, and simplest, it is likely the case that forbearance is not permanent – should the Commission find at a later date that the conditions giving rise to forbearance have changed (including simply that it is, in the (current or future) Commission's view, in the public interest to discontinue its forbearance) any provision of Title II could come back into force. Thus the argument that forbearance pares back the most onerous provisions of Title II, thereby easing concerns about the burdens Title II regulation imposes and the expansion in scope of the FCC's power, is illusory. These problems are only avoided by the grace of the Commission's own beneficence. Forbearance, in other words, is no cure to the previous concerns expressed by the Court in *UARG* – concerns which Justice Scalia explained provided sufficient basis on their own (*i.e.*, independent of the EPA's tailoring efforts) to reject the EPA's rules.

Second, and more nuanced: the FCC's forbearance power is limited, such that forbearance can only be granted under certain circumstances. This means that forbearance may not be sufficient to address concerns of excessive burden and statutory structure. Section 10 requires that the Commission forbear from enforcing provisions of the Communications Act subject to three conjunctive conditions: enforcement is not necessary to ensure just and reasonable prices and practices, enforcement is not necessary for the protection of consumers, and forbearing from

[46] Section 10 of the Communications Act of 1934, as amended, codified at 47 USC 160.

enforcement is in the public interest.[47] None of these factors considers whether forbearance is necessary to make the Commission's preferred construction of the statute work in practice; and the fact that forbearance may be necessary in order to make the statute work is not on its own sufficient to trigger forbearance. As a result, it is entirely possible that the Commission's classification of broadband Internet access under Title II could yield burdensome, or otherwise problematic, results sufficient to render the classification impermissible.

Chairman Wheeler's proposed "modernization" of Title II is therefore based on authority just as weak as that upon which the EPA based its "tailoring rule." The FCC does have authority to forbear from enforcing portions of the Communications Act – and the EPA, like every agency, has some discretion to determine how it exercises its authority. But the FCC's power is not discretionary: Section 10 is written in the imperative requiring that "the Commission shall forbear" if certain conditions are met. The statutory design therefore serves to *constrain* the Commission's discretion over what provisions of the Communications Act it enforces. It is not a mechanism to allow the Commission to "rewrite clear statutory terms to suit its own sense of how the statute should operate."[48] It allows the Commission to trim the fat of the statute, not to excise tumors threatening to consume it. As such, like any other agency, the FCC cannot "adopt . . . unreasonable interpretations of statutory provisions and then edit other statutory provisions to mitigate the unreasonableness."

The Chairman's approach is made even more egregious because it is not compelled. The DC Circuit made clear – at a time prior to the Court's decision in *UARG* – that Section 706 presented a reasonable means to the Chairman's desired ends. Had the Chairman opted to follow the path set forth by the DC Circuit, we would not find ourselves in the present situation, and the Commission would be on much firmer legal grounds. Rather, the Chairman has elected to pursue the more aggressive path. The dangers of this discretionary path are substantial. As explained in the EPA context, "Since, as we hold above, the statute does not compel EPA's interpretation, it would be patently unreasonable – not to say outrageous – for EPA to insist on seizing expansive power that it admits the statute is not designed to grant." The Chairman's chosen path is similarly unreasonable – and equally outrageous.

This is particularly the case given how little sense the Chairman's explanation for his sudden embrace of Title II makes. He has explained that at some point (roughly contemporaneous with the President urging him to take a Title II-based approach), he "became concerned that the relatively untested [Section 706-based] 'commercially reasonable' standard might be subsequently interpreted to mean what was reasonable for the ISP's commercial arrangements." In his words, that "was a possibility that was unacceptable," and this led him to embrace Title II. This explanation, however, simply make no sense. First, the DC Circuit had made clear that Section 706, at least in its view, provides the Commission with the necessary authority. And second, the concern that the "commercially reasonable" standard proposed in the May NPRM would be interpreted to protecting ISP's commercial interests could be trivially addressed with a

[47] *See* 47 USC 160(a). The statement of the first requirement is a paraphrase.
[48] *UARG* at 2446 ("We reaffirm the core administrative-law principle that an agency may not rewrite clear statutory terms to suit its own sense of how the statute should operate.").

single sentence, along the lines of "Determinations of 'commercial reasonableness' shall give substantial weight to the effects of those agreements upon consumers, and no agreement shall be deemed commercially reasonable that is not also substantially in the consumer interest."

Substance of rules under Title II

Finally, just as with rules adopted under Section 706, there is reason to expect that any rules adopted under Title II will be rejected as arbitrary and capricious. As explained previously, an outright ban on paid prioritization, in particular, runs afoul of the overwhelming weight of evidence in the economic and technical literature. This literature consistently shows that paid prioritization can be either beneficial or harmful to consumers – the specific effects in any case are fact-specific. Any ban on paid prioritization necessarily forecloses potential consumer benefits in exchange for uncertain gains.

Counterintuitively, the case for such a ban may be even weaker under Title II than under Section 706. This is because Section 202 prohibits only "unjust or unreasonable discrimination in charges." Indeed, Chairman Wheeler has testified before Congress that Title II does not prohibit paid prioritization.[49] The statutory language – "unjust or unreasonable" – must be given meaning, and the overwhelming weight of the academic literature makes clear that paid prioritization can in many cases be both just *and* reasonable. Given the weight of the relevant literature showing that such treatment can be pro-consumer – especially without any serious evidence of actual consumer harm resulting from prioritization, paid or otherwise – it is not hard to expect the courts to reject the Commission's substantive rules both as lacking and contrary to factual support.

Conclusion

The FCC has embarked on one of the most aggressive expansions of authority of any agency in our nation's history. Clothing its efforts in the statutory ambiguity of the Communications Act's dated and circular definitions and taking comfort in dicta from *Brand-X* does not change this fact. As Chairman Wheeler has made very clear, the Internet is different in kind – in scale, scope, and nature – from what the Commission was granted authority to regulate. It really is "the most powerful network in the history of mankind," central to our modern economy. And, in giving the Commission authority in 1934 to regulate the telephone monopoly, and again in 1996 to unwind and eventually deregulate that monopoly, Congress did not give the Commission authority to regulate the Internet.

In this *Perspectives*, I have sought to map out parts of the road that lies ahead. It is, of course, entirely possible that the arguments I highlight will fail in practice – but, I believe, they also have a good chance of carrying the day. This means that we can say, with great certainty, that the road ahead is uncertain for the Commission and the industry. And, most tragically, it will be rocky for consumers – consumers who have already been subjected to years of discord as net neutrality

[49] House Energy and Commerce Committee Hearing on FCC Oversight (May 20, 2014), *recording available at* http://www.c-span.org/video/?319457-1/fcc-chair-testifies-net-neutrality ("As you know, Title II, there is nothing in Title II that prohibits paid prioritization. As a matter of fact, we have all kinds of paid prioritization...").

proponents have sought sweeping new power for the FCC to regulate the Internet. It is all the more a pity given the many opportunities the Commission has had (and, indeed, still has) to avoid the whole mess. But instead: It almost certainly looks like "hi ho hi ho it's off to court we go!"

* Justin (Gus) Hurwitz, a member of the Free State Foundation's Board of Academic Advisors, is an Assistant Professor of Law at the University of Nebraska College of Law.

The Free State Foundation is an independent**,** nonpartisan free market-oriented think tank located in Rockville, Maryland.

THE FREE STATE FOUNDATION
A Free Market Think Tank for Maryland.....Because Ideas Matter

Perspectives from FSF Scholars
September 30, 2013
Vol. 8, No. 25

Two Sides of the Internet's Two-Sidedness: A Consumer Welfare Perspective

by

Justin (Gus) Hurwitz *

This month's Open Internet argument before the DC Circuit was exciting for many reasons – including many aside from the substantive arguments about the FCC's authority to issue the rules under review. I'd like to look closely about one of these reasons in particular: the argument that consumer broadband Internet providers are part of a so-called two-sided market, and the implications that this argument has for consumers. Whether or not this is a two-sided market (or could be, absent the FCC rules) is important to understanding how the Open Internet rules affect the development of the broadband Internet market – and ultimately how they affect consumers.

But first, what is a two-sided market? The most basic definition of a two-sided market is a market with two distinct groups of consumers for some good, where the number of consumers from each group consuming that good affects the demand of the consumers of the other group for that good. Examples should help: nightclubs are generally two-sided markets, because the number of men at a given club affects how much women will want to go to that club, and vice versa; health insurance plans are two sided markets, because the number of doctors in a plan's network affects consumers' willingness to be part of that plan, and vice versa; computer operating systems are another example, because the number of applications available for an operating system affects consumers' willingness

to use that operating system, and the number of potential users affects programmers' willingness to develop applications for a given operating system.

So do broadband Internet service providers operate in a two-sided market? It seems so: all else equal, consumers are going to have greater demand for an ISP that gives them access to a wider range of content providers (such as Google and Netflix) rather than a smaller range. And content providers are going to care more about connecting to backbones and broadband ISPs (the largest of which are tier 1 and tier 2 backbone providers) with more customers rather than fewer.

But there is something missing from the discussion so far: why do we care if a market is two-sided? Because in most two-sided markets, the purveyor of the intermediary goods that the two sides are consuming – that is, the owner of the nightclub, the HMO provider, the OS developer, or the broadband ISP – sets different prices for each side of the market in order to maximize the value of the market. Why? Because doing so allows it to increase its own revenue! So nightclubs let women in for free, to attract more men; HMOs might make it easy for doctors to join in order to attract more customers; and OS developers may give away their APIs and SDKs for free, to make it easier for programmers to write the programs that will attract users.

This is part of the argument that Verizon made before the DC Circuit: the Open Internet rules, by preventing Verizon from charging firms like Google and Netflix for access to its network, prevent this market from behaving like a two-sided market. (As a brief technical aside, it is entirely possible for a "two-sided" market to have more than two sides. For example, there could well be four "sides": consumers, high-value high-bandwidth services (Netflix), high-value lower-bandwidth services (Google, Facebook), and all other services. The economic analysis of such "multi-sided" markets is more complicated than of two-sided markets, but the idea is the same.)

The poignancy of this argument is that the economic literature makes amply clear that different price structures have powerful effects on the value of a given two-sided market to consumers. A nightclub that subsidizes women's entry may well attract far more consumers (both men and women!) than one that charges men and women the same cover or than one that subsidizes men's entry. The Open Internet rules impose a given price structure on the broadband market, Verizon's argument goes, without any evidence that that price structure is, in fact, the one that maximizes the value that that market creates for consumers.

Verizon's basic argument here is almost certainly correct: figuring out the best price structure in two-sided markets is complicated, and there is little reason to believe *a priori* that the Open Internet rules' prohibition on charging content providers is optimal. To the contrary, the economic literature suggests that the Open Internet rules can have a negative effect on the value created by the Internet, and that allowing broadband ISPs to charge content providers can benefit consumers and increase infrastructure investment.

What is more, the literature also suggests that the Open Internet rules are most likely to be harmful where there is little competition between ISPs – in other words, the world in which Network Neutrality advocates are generally eager to remind us we live in, where most consumers have access to only a few ISPs, is the world in which the Open Internet rules are most likely to be harmful.

The takeaway is that today's broadband Internet market is precisely the sort of market in which the FCC's "prophylactic" approach is inappropriate. It is, instead, one in which we should seek out opportunities to experiment with multisided price structure – and even reward firms for taking the risk of experimenting – in order to maximize the value of the Internet to consumers.

But wait! There's more! Many strands of the literature on two-sided markets introduce an additional requirement for a market to be two-sided: the parties on either side of the market need to be unable to bargain around the price structure. This means, to take a pointed example, that if Netflix can pass any charges that Time Warner Cable levies upon it back onto its customers, then Netflix is not party to a two-sided market.

This example offers another motivation for content providers – at least, those able to pass costs back on to their customers – to favor Open Internet rules. If a content provider's services stress an ISP's network, that ISP faces two options: either allow the quality of its network to fall or invest in upgrades. If it incurs the cost of upgrading the network, those costs will be passed on to its customers. In either case, the ultimate cost is borne by consumers.

This allocation of costs – spread as it is across all of the ISP's customers – makes little sense. Far preferable, at least for consumers and ISPs, would be a price structure that initially allocates the cost of network upgrades to those whose use of specific services requires those upgrades be made. Otherwise the least resource-consuming users are subsidizing the most resource-consuming users. And this creates an incentive for all users to consume more resources (the low-use users because they're paying for it, and the high-use users because their use is being subsidized). (Query, for readers interested in game theory: where's the stable equilibrium?)

Of course, this system isn't bad for everyone: the content providers prefer the Open Internet model. Under this model, they are shielded from needing to increase their prices – that consumer-angering move is passed on to the ISPs. They also benefit from the subsidy from low-use users to high-use users, because this effectively subsidizes use of the content providers' services. What's more, they benefit yet again from the incentive that this system creates for users to consume more resources – to consume, that is, more of the content providers' services.

It should be noted that this is exactly the sort of feedback loop that we expect to see in a two-sided market – indeed it is the mechanism that causes these markets to create value for consumers. But this is entirely consistent with Verizon's and other ISPs' presumed desire to experiment with two-sided price structures. It could well be the case that the

value of this market is maximized by ISPs, instead of content providers, passing infrastructure costs on to the consumers (query how they would be apportioned, evenly across all consumers or by usage). But if that is, indeed, the value maximizing outcome, it could obtain by agreements between ISPs and content providers, whereby they split the increased revenue that results from increased consumer demand.

Either way you have it – whether or not broadband Internet service is a two-sided market – the Open Internet rules are potentially a bad deal for consumers. We can't say categorically that this is the case. But "I dunno, maybe?" isn't a good enough basis for policy. What is clear is that these rules are a subsidy to content providers.

It's a shame that broadband Internet services and content providers are the two sides of these policy debates, when the consumers are the only side that really matters.

* Justin (Gus) Hurwitz, a member of the Free State Foundation's Board of Academic Advisors, is an Assistant Professor of Law at the University of Nebraska College of Law.

The Free State Foundation is an independent**,** nonpartisan free market-oriented think tank located in Rockville, Maryland.

What the Open Internet NPRM is – and isn't – about

by: Gus Hurwitz

May 12, 2014 6:00 am

COMMUNICATIONS

Internet by Shutterstock

tags: net neutrality, Title II, Wheeler

We're three days away from the FCC meeting at which the commission is scheduled to vote on the chairman's Notice of Proposed Rulemaking to revise the Open Internet rules. I want to take this post to discuss what this NPRM is – and isn't – about.

Let's start with what the NPRM isn't about. It isn't about the future of the Internet. It is not about whether the FCC will or will not regulate business practices on the Internet. It is not about whether the FCC will or will not allow "fast lanes." Anyone who tells you otherwise is advocating for a preferred outcome. In many ways that's fine: most people interested in this NPRM have preferred outcomes, and many are so steeped in the issues (and their beliefs about them) that the "future of the Internet" rhetoric is useful shorthand. But for most folks — those who have never heard of an NOI, NPRM, FNPRM, R&O, or the APA — the shorthand is at best confusing.

Today, the FCC has broad authority to regulate the Internet. Under Section 706, if a firm were to implement a consumer-harming business model that violated the principles of network neutrality, the commission could bring an administrative action against that firm: the commission could sanction the firm and require it to forego the business practice. Indeed, under the generous deference afforded agencies such as the FCC, the commission could almost certainly take such action if it were presented with evidence that such business plans only harmed consumers indirectly, for instance by adversely affecting incentives for edge providers to innovate or invest in new technologies.

In other words, the "rules" that net neutrality advocates are demanding already exist. The Communications

Act already gives the FCC power to investigate and take action against business practices that violate network neutrality, and instruct it to use that power when doing so is in the public interest, convenience, or necessity.

The Open Internet NPRM isn't about anything that has grabbed the attention of the media or protesting public. The NPRM is about the decidedly unsexy legal question of how to implement the already existing statutory authority. Should the FCC implement clear ex ante rules that may harm consumers by foreclosing the development of pro-consumer business practices? Or should it instead rely on general standards that are enforced on a case-by-case basis as bad conduct arises? Importantly, in either case the FCC would need to bring an administrative action against a firm that violated neutrality principles. If the Commission implements broad ex ante rules, it risks that they will be invalidated by the courts – for a third time! – when they try to enforce them. If, instead, the commission relies on standards, the commission will have a slightly harder time challenging problematic conduct, but will face less risk that the entire Open Internet regime will be invalidated.

So the question of whether to implement clear rules or rely on case-by-case adjudication isn't about whether to allow or proscribe "fast lanes." It's about how the FCC will take action against firms that harm consumers, the burden of proof that it will face when it takes such action, and the risks that it faces should it lose.

This is the sort of issue that only a lawyer can love – and even then it's one for which many lawyers have little patience. Every law student learns during their first year (usually their first semester) about the blurry distinction between rules, standards, and principles. And then they spend the remainder of their time in law school struggling with this distinction as it plays out in various contexts. There is always a natural preference for rules, which provide clarity and the comfort of certainty; but they are also always under- and over-inclusive. Standards provide greater flexibility so to reduce the costs of this under- and over-inclusiveness, but they come at the cost of less certainty and greater costs when it comes time to enforce them. That is the entire debate over the Open Internet NPRM in a nutshell.

The same is basically true for arguments over reclassifying the Internet as a Title II ("common carrier") service. Contrary to commonly held opinion, Title II services don't prohibit discrimination. Title II isn't a magic wand that requires common carriers to be neutral. It only bars "unjust and unreasonable discrimination." Given the overwhelming body of technical and economic literature showing that discrimination can benefit consumers and facilitate new services — and that the Internet never has been "neutral" in the sense that most net neutrality advocates use the term — there is no reason to believe that reclassification would prevent firms from adopting business models to which net neutrality advocates object. Indeed, think about any services traditionally understood as common carriers: trains, roadways, parcel and post, electricity and other utilities, and even communications. Each of these services has always allowed for differential levels of service — and, when a customer of any of these services gets to be large enough, common carriage rules have always allowed individualized and preferential services to be negotiated in one form or another.

In other words, as with the rules vs. standards debate in the context of Section 706, the debate here is about the relative certainty afforded by reclassifying versus not reclassifying, and the relative costs of

taking action against firms under either regime. Whether or not the FCC reclassifies, it would still need to take action against a firm engaging in potentially harmful conduct – and the arguments against and in favor of the firm's conduct would be roughly the same under either regime. Under Title II, the FCC would likely have an easier time winning a case against a bad-acting firm. But the cost is that it would be harder for firms to develop new, pro-consumer, business models. Doing so would require receiving FCC permission under a process known as forbearance – a process that typically takes 12 to 15 months (or, in Internet-time, forever).

Finally, this isn't the FCC's only bite at the apple. The commission is free to change its rules, almost at a whim. So, should the commission adopt rules this year that prove to be too lax, it can change those rules next year (or even later this year!). Importantly, if the commission adopts rules that are too lax, it will be relatively easy to tell (firms will adopt harmful policies), and it will be easy for the commission to change the rules in response. Indeed, having direct evidence that stronger rules are needed, it will be easier to adopt tougher rules than it is today. On the other hand, if the commission adopts too tough rules, we are unlikely to ever see the harm – it's hard to see when something good *doesn't* happen. So, if the commission adopts too strong rules today, we're probably stuck with them. We'll all be poorer for it, though we may never know.

It's very easy to get caught up in the passion and rhetoric surrounding network neutrality. But this issue before us right now – and before the commission on the 15th – isn't about network neutrality. It's not about whether the commission has the power to take action against firms that violate network neutrality principles. The commission already has that authority. The issue before us right now is the boring, technical, lawyerly question of how to implement that power. Should we rely on rules or standards; what are the relative costs of each; what are the relative risks of each; how do they affect burdens of proof.

My recommendation to the chairman, should he be reading this in a moment's refuge from the angry mob outside his office: if the commission does not adopt the NPRM on Thursday, let the issue die there. You already have authority to take action against problematic conduct. The lack of a rule-making will make it harder to take such action (especially if the commission were to seek fines or other due process–protected sanctions) — but fundamentally, you still have the tools to protect the public interest. And you have more important things to do than tend to the care and feeding of your predecessors' albatross.

An unfounded principle: Ammori's non-neutral network history

by: Gus Hurwitz

November 13, 2013 6:00 am

COMMUNICATIONS INTERNET

NetworkDesign by Shutterstock

tags: net neutrality

Marvin Ammori's recent Wired article on the DC Circuit's likely rejection of the FCC's Open Internet Order continues to garner attention. Richard Bennett, our newest colleague here at TechPolicyDaily.com, wrote a nice takedown of Ammori's article last week. I won't belabor the points that he made here.

But I do want to correct one very important mistake in Ammori's piece. Ammori calls net neutrality "the foundational principle" of the Internet. Not just "a" foundational principle (which it isn't), but "the" foundational principle. Every time I have read his article over the past week, this statement makes my skin crawl. It is just wrong – flat out, plain and simple, no-basis-in-reality wrong.

Let's start by looking at the actual "foundational principles" of the Internet – to the extent that there are any. The Internet resulted from two things: a new need and a new technology that was able to address that need. The history as I describe it is corroborated in many sources. A good starting point – one authored from a perspective generally sympathetic to Ammori's views – is the Internet Society's Brief History of

theInternet.

In the 1960s, researchers Paul Baran, Donald Davies, and Larry Roberts began working on a new technology – packet switching – to enable computer-to-computer communications. (Their work was occurring at the same time as similar work by others at other institutions – the idea of packet switching was "in the air" at the time.) This technology was highly theoretical, but held the promise of facilitating multi-user access to time-shared computers in ways that traditional circuit-switched technology could not accommodate. ARPANET, designed by Roberts, funded by DARPA and built by BB&N, began operating in 1969.

After ARPANET had been running for a few years, Bob Kahn of DARPA developed an interest in connecting it with two other packet-switched research networks (one satellite-based the other radio-based) and "internetting" (as Kahn called it) was born. One of the many remarkable things about packet switching is that data traversing any network could be broken into packets, which could then be transmitted across another network and reassembled on yet other networks; and once data was put into packets, it could traverse any network, independent of that network's design. The Internet grew out of ARPANET, which was developed to facilitate packet-based internetworking.

As told in the Internet Society's account, these ideas gave rise to the "key underlying technical idea" of "open-architecture networking" – the closest thing to a "foundational principal." In this model of networking, "the individual networks may be separately designed and developed … . Each network can be designed in accordance with the specific environment and user requirements of that network. There are generally no constraints on the types of network that can be included … , although certain pragmatic considerations will dictate what makes sense to offer."

This "key underlying technical idea" doesn't sound anything at all like Ammori's "key foundational principle." To the contrary, it sounds as though networks *can* have different technical characteristics, which may result in different treatment of data.

This packet-based approach to internetworking gave rise to several other considerations. For instance, in order for the internetworking concept to work, hosts on any one network had to be able to connect to hosts on any other network – this, after all, is the basic purpose of internetworking. This is part of what is referred to today as the end-to-end principle. And, the packet-based nature of the network – the idea that packets could traverse any sort of network – means that reliable delivery of packets cannot be assumed. Rather, the best the network can be assumed to offer is its "best effort" to deliver packets. But these considerations result from the technological nature and purpose of the Internet – they are not themselves key design principles. (Other aspects of the end-to-end and best effort principles, and related ideas such as layering, also fall out of the packet-based approach to networking. These ideas, which are beyond the scope of this essay to discuss, have frequently been mangled to support network-neutrality- and open-access-like ideas. Fundamentally, however, they, too, are based on a pragmatic approach to engineering, and are not meant to be taken as absolute principles. See, e.g., RFC 3439 ("Layering considered harmful"); and generally Richard Bennett,Designed for Change.

This is a key point about the Internet. It was not built around high-falutin' principles or ideals. It was not designed to be "generative," or to promote innovation. It was, itself, an exercise in theory (packet switching) and innovation as part of a pragmatic attempt to solve complicated engineering problems. Any "principles" were identified after the fact, to explain what had worked. Such principles demonstrate only what was sufficient – not what was necessary – for the Internet's success.

More important, even if there were key principles, they do not inevitably lead to a "network neutrality" principle. Nothing requires "neutral" treatment of packets – there was a practical preference early on for simple routers, to avoid technological complexity. But this was a pragmatic consideration, not a technical one. Similarly, decisions over which services to offer (which could certainly include which networks to interconnect with) could be subject to pragmatic considerations – both technical and economic.

These latter comments lack direct evidence to support them. This is because "network neutrality" is a new concept. No one was discussing it in ways relevant to modern usage in the Internet's formative years. There was little need to discuss it then, largely because routing technology wasn't sophisticated enough to enable differential treatment of packets, and there was little technological or economic need for such technology.

Although contemporaneous discussion of "network neutrality" may be lacking, we are not without evidence about the permissibility of differential treatment of data.

For instance, what we now call TCP/IP was originally only TCP. In the 1970s, engineers recognized that not all of the features implemented in TCP were needed for all types of applications – indeed, some of these features could be harmful to application such as streaming voice. So TCP was split into separate network (IP) and transport (TCP) functions, and a new transport protocol (UDP) added. This was an express recognition that some applications may require differential treatment.

Some may object that the division of TCP into TCP/IP and UDP/IP only affected the transport layer, moving services that did not belong in the network layer to a separate layer – thereby increasing neutrality. This argument does not hold up to scrutiny, however. During this same period, researchers also worked on a new protocol – the Internet Stream Protocol (See IEN 119 and RFC 1190) – which is an IP-layer (network-layer) protocol to support streaming media alongside TCP/IP and UDP/IP.

There is also work dating back to the at least the early 1980s to implement network-level Quality and Type of Service features. These features would allow router-based differential treatment of packets based upon, for instance, the technical requirements of the application (low latency vs. high bandwidth) or economic considerations (transit or termination costs of one network vs. another).

The idea of network neutrality is a thoroughly modern concept, cut from whole cloth by those who believe that the historical fact of a mostly (but not entirely!) neutral Internet is responsible for the success and growth of the Internet. But the engineers who developed the Internet were fiercely pragmatic, and neutrality was a consequence of that pragmatism: deliberate non-neutral treatment of packets was prohibitively complex for the technology of the time, development of such technology a distraction from addressing the needs of the network, and it didn't address any technological, social, or economic goals.

We should not take the fact that neutrality was the norm to mean that it was a necessary condition for the Internet's growth and success. Correlation between neutrality and the Internet's growth and success does not equal causation. Those who would insist upon a neutrality mandate bear the burden of proving its wisdom – and surely that burden is not met by pointing to a fictional foundational principle.

Five faulty premises, part 4: Innovation requires open access and an open Internet

by: Gus Hurwitz

February 21, 2014 6:00 am

COMMUNICATIONS

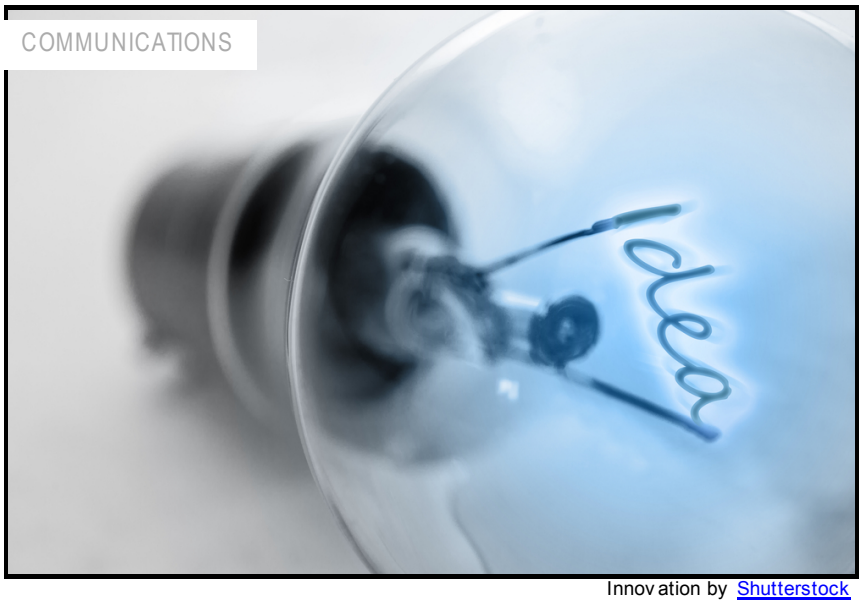

Innovation by Shutterstock

tags: faulty premises, FCC, innovation, net neutrality

In this installment in a series of blogs based on my recentFree State Foundation Perspectives article, I look at the fourth of five faulty premises, that innovation requires open access, and in particular a (so-called) open Internet. While this isn't directly responsive to FCC Chairman Wheeler's announcement yesterday of his plans to initiate a new Open Internet rulemaking, given this faulty premise's relevance to the subject I offer some comments on such a rulemaking at the end.

This faulty premise is a doozy. It is the beautiful premise that launched a thousand ships on the sea of Network Neutrality. But its beauty is skin deep. While it is true that open access can facilitate some types of innovation, it both precludes other forms of innovation and imposes costs of its own. In the telecommunications context, open access is mostly about network neutrality – the idea that broadband providers should not be able to charge users or content providers for preferential access to specific services, let alone block specific content or services entirely (absent some compelling legal or technical justification).

It is unquestionably the case that open access can facilitate certain types of innovation. It reduces R&D and

other transaction costs (especially search and negotiation costs to get permission or access to use existing infrastructure) and reduces opportunities for rent extraction by those who otherwise control an infrastructure. On the other hand, it makes some forms of innovation more expensive or difficult to implement. There are substantial literatures showing the benefits of vertical integration and the importance of defining proper modular boundaries. Nowadays, however, this point can be made more simply: Apple's hardware and software designs are part of a tightly-controlled, vertically integrated, closed product ecosystem. Apple would not exist if we had the equivalent of network neutrality for computer hardware or software. This does not mean that either an open or a closed model is necessarily better in any given case; it does mean that we want a more nuanced approach than one that mandates either approach in every situation.

The scale is tipped even further against mandated open access in the case of the Internet. This is because the Internet is a two-sided market – a market in which two or more distinct groups of consumers are brought together via some intermediary platform. That is, users and Internet content providers (e.g., firms like Google, Facebook, and Netflix) reach each other via the Internet. I looked at this issue in some depth in a prior FSF Perspectives piece; indeed, I first explored some of the early literature on multi-sided markets and the implications multi-sidedness has on the Network Neutrality debate in a 2006 article.

The crux of the two-sided markets analysis is that the platform that brings the different sides together – that is, broadband Internet access providers – ordinarily charge either or both sides of the market for access to the other. How much to charge each side, including whether to charge either side nothing or even to subsidize one side's access to the platform, involves a complex set of tradeoffs – and, most important, how much each side is charged can have substantial effects on the social value of the network. Critically, and I will say this in italics because it is so important, *the literature studying two-sided markets consistently shows that there is no reason to believe that a network neutrality rule benefits consumers, and consistently shows that such a rule can harm consumers*.

A network neutrality rule is really little more than a subsidy from the consumer side of the market to the content provider side of the market. Some, but not all, content providers benefit from this rule. Other content providers may be harmed by such a rule – especially those who offer, or would like to develop, services that would benefit from enhanced quality of service features or other features that may require some integration with Internet service providers.

Even more problematic, a network neutrality rule can harm consumers. It prevents ISPs and content providers from working together to offer innovative new products that consumers want. More tragic, it prevents these providers from developing lower-cost service packages – packages that could expand opportunities for access to currently underserved and disadvantaged communities. These rules likely increase cost of access and limit the development of potentially cheaper offerings that are more responsive to consumer demands – this is exactly the opposite of good telecom policy.

This point relates back to a concern in the first faulty premise considered in this series: the paramount importance of respecting consumer preferences, and not substituting the Washington-Silicon Valley-Ivory Tower views of what consumers should want for what they actually do want (and, more importantly, need).

By requiring that every consumer's Internet connection offers full-fare, first-class service, complete with movies, television, and free drink service, we price consumers who would be happy with discounted-fare economy Internet service out of the market.

I don't mean to give away the barn. The key takeaways from the literature in this field are all nuanced – different price structures "can" or "may" benefit or harm consumers. In some cases, "non-neutral" price structures may benefit consumers, in some it may harm them, and conversely. (Noted paraleptically, my own reading of the literature suggests that, given current market structures, non-neutral pricing is likely to be better for consumers than neutral pricing.) But this does not mean that we should prescribe *ex ante* prophylactic pricing rules – rather, it means that we should monitor conduct and pricing in the Internet ecosystem and be ready to bring *ex post* actions against pricing decisions that are demonstrably harmful to consumers.

Some additional comments bear mention in light of FCC Chairman Wheeler's announcement yesterday that a renewed Open Internet rulemaking process is in the offing. The goal of those rules must be to protect consumers, not to protect content companies. What's good for Google is not necessarily good for America.

Whatever rules the Commission may ultimately adopt, the Commission should be careful that they do not proscribe pro-consumer conduct. Given the difficulty of knowing ex ante whether any specific conduct is likely to benefit or harm consumers, whatever rules the Commission ultimately adopts likely should be limited to general principles – they should not define conduct that is to be prospectively permitted or prohibited, but rather (at most) indicate certain types of conduct that may bear scrutiny from the agency and the terms under which that conduct will be evaluated. Should the Commission take such an approach, it may yet craft an approach that passes judicial muster, provides useful guidance to agencies, and – most important – protects consumers from harmful conduct on the one hand while allowing them to benefit from pro-consumer innovation on the other.

Markets, regulation, autonomy, and dignity

by: Gus Hurwitz

March 14, 2014 6:00 am

COMMUNICATIONS

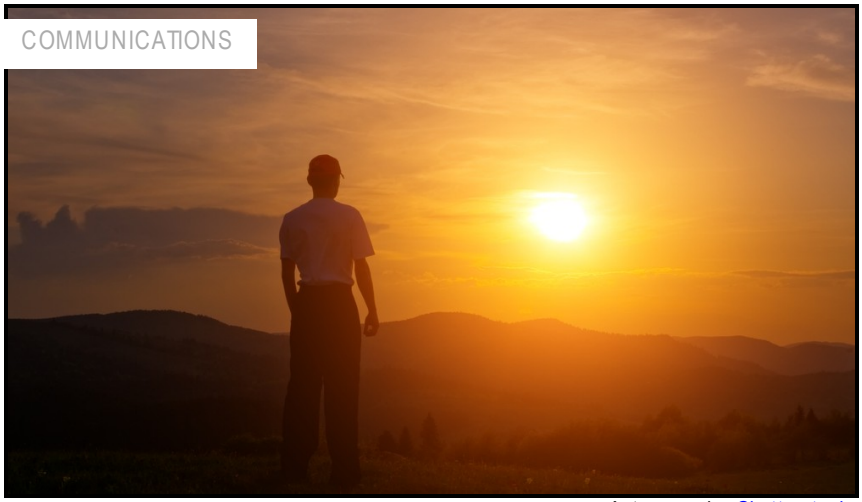

Autonomy by Shutterstock

tags: Comcast, M&A, Netflix

Following news of the proposed Comcast/Time Warner Cable merger and the Comcast/Netflix interconnection agreement, the tone of telecom and tech policy discussions has fallen precipitously. Most commentary about these deals has been immediate and reflexively hostile – based in emotion and fear rather than in facts. This response is frustrating to watch – though at some level it has been a return to normal, following a few months of atypically reasonable discussion leading up to the DC Circuit's Network Neutrality decision.

At its core, the response to these deals reflects popular anti-market, pro-regulatory, sentiment. I would like to spend some time thinking about these attitudes. Perhaps surprisingly, I find them quite sympathetic – the pro-regulatory view is intuitively appealing and at some level a reasonable one to hold. While I disagree with it, it is important to understand why this view is so durable.

The individual is central to competing views about regulation and the market. Market-oriented advocates generally view the market as respecting and promoting individual autonomy and dignity – and they fear that regulation robs individuals of this dignity, denying them their basic freedoms of association and conscience. Advocates for stronger regulation generally have the opposite perspective, distrusting markets and fearing that they treat individuals as commodities, robbing them of same dignity that market-oriented thinkers seek to protect. The pro-regulatory view sees regulation as a way preserve individuals' dignity against a hostile market.

Which view is right? The answer, of course, is both – for reasons that I believe are particularly important for free-market advocates to understand. If the free-market view is to be internally consistent, its advocates need not only to focus on why they are right, but also to respect why its detractors are not wrong. The discussion below proceeds in four parts: an explanation of the market-oriented view; discussion of what that view overlooks, and why this supports the idea that markets do not support individual autonomy; the economist's response to this accusation; and the need for market-oriented thinkers to be responsive to these concerns.

Part I: The market-oriented view – Complexity and the mysterious invisible hand

I fall into the school of economic thought that emphasizes economics as the science of self-interest and mutually beneficial exchange. Broadly speaking, individuals make decisions that make themselves better off – this means that any agreement between multiple parties must make each party to the agreement better off. Certainly there are cases where this isn't true – most classically in cases of monopoly and diffuse negative externalities (e.g., pollution). But by and large, arguing from both theory and empirics, the moral case for free-market economies is strong. (See, for instance, Alex Tabarrok's discussion of economists as leading opponents of slavery, Mill's classic argument for equality between the sexes, and Deirdre McClosksy's recent exploration of the question "Are Market's Moral?")

Markets certainly are not perfect. No market is perfectly efficient; few markets come anywhere near the ideal of perfect competition. It is the great failure of contemporary undergraduate economics education that most students come away believing that perfect competition is an assumption of economics. (Indeed, the first task of graduate-level economics education is to break down the undergraduate-level caricatured understanding of economics.)

But economics educators use this approach for a reason: the operation of economics is confoundingly complicated. As Hayek famously explains, "The curious task of economics is to demonstrate to men how little they really know about what they imagine they can design." The challenge of this task is that the functioning of markets is overwhelmingly complex, and the apparent problems with them are so patently obvious. Anyone who has ever felt they have paid too much for some good or service has identified what they view to be a problem with the market (the market's price was too high!), and the solution to this problem is maddeningly obvious (the government should require lower prices). It generally takes years of graduate level study to develop an intuitive sense that interactions between billions of individuals gives rise to a coherent and beneficial market – one guided by Smith's invisible hand – and to overcome the natural tendency to view so many market-based interactions as unfair.

Part II: What the market-oriented view forgets

This is where we find that the motivation for the regulatory approach is also based in concerns for individual autonomy, respect, and dignity.

It is often forgotten that, on the same page as he invokes the invisible hand, Smith also notes that "it is the maxim of every prudent [individual] never to attempt to make at home what it will cost him more to make than to buy." This is the basic mechanism animating the invisible hand – it is the same as my initial

explanation of mutually beneficial exchange. Individuals only enter into agreements that make them better off as compared to their next best alternative.

But there is an assumption built into Smith's prudent individual's "make or buy" decision: that individual needs to know how much it will cost to make a good compared to buying it. An initial assessment of his ability to do this underlines the pro-regulatory concern that markets rob individuals of their autonomy: few people know how much it would cost to build a telecommunications network; and few trust that the market prices for telecommunications service are fair. And it is perfectly reasonable that they don't trust the market prices: just as they don't know how much it would cost to construct their own network, they lack a baseline price against which to compare the market price. Thus, they are forced to accept a price that they do not trust because they lack – and cannot reasonably be expected to have – the facts needed to make an informed decision.

This view is captured in Brian Fung's recent article on the Comcast-Netflix deal: "The central problem here is a lack of information. We don't know how Comcast's deal with Netflix is structured and how much money is changing hands. That makes it hard to judge whether critics are right when they claim Comcast is engaging in anticompetitive business practices. And it makes it hard to know whether consumers will wind up paying more as a result of the deal. We could fix this by having the government require Comcast to be more forthcoming about its Netflix relationship."

Part III: The economist strike back

The economist's response to this is that the relevant parties – Comcast and Netflix – have excellent understandings of the underlying "make or buy" decision. Netflix had been self-provisioning its content distribution (i.e., "making" it), but had the opportunity to enter a better deal with Comcast to buy content distribution. There may be a possibility that the parties have entered a deal that is harmful to consumers – but this is why we have the FCC and DOJ. It is almost a foregone conclusion that both agencies had copies of the agreement between Comcast and Netflix within hours of the deal being announced.

More generally, the economist's response to the concern that Smith's "make or buy" decision forces individuals to choose between something they don't understand and something they don't trust is that both questions are answered by the market. The market price for a given good or service provides a first approximation of the cost of making that good or service. This is true even if the vast majority of market participants don't know enough to make the make or buy decision. So long as some (relatively small) portion of the market has a sense of these costs, competition – that is, everyone else in the market, each with their own specialized knowledge, making the same make-or-buy decision – will ensure that market prices reflect cost.

The economist will also say that we must look at the other side of the equation. Just as the individual must choose between making and buying, society must choose between competition and regulation – and must recognize that neither is costless. Just as there is no such thing as perfect competition, there is no such thing as costless or perfect regulation. To the contrary, the history of regulation is fraught with failure. As discussed by Ronald Coase in a 1997interview:

> When I was editor of The Journal of Law and Economics, we published a whole series of studies of regulation and its effects. Almost all the studies – perhaps all the studies – suggested that the results of regulation had been bad, that the prices were higher, that the product was worse adapted to the needs of consumers, than it otherwise would have been. … I can't remember one that's good. Regulation of transport, regulation of agriculture – agriculture is a, zoning is z. You know, you go from a to z, they are all bad. There were so many studies, and the result was quite universal: The effects were bad.

Part IV: What must economists do?

The challenge that we – those of us who believe in the economic and free-market approach to regulation – face is that laid out by Smith, Hayek, and Coase. Our task is to explain on the one hand that despite its near-impenetrable complexities, the market tends to present individuals with a reasonable "make or buy" decision, and to explain on the other hand that despite regulation's self-evident simplicity its actual practice is fraught with failure. The market, though it often does not *feel* this way, treats individuals with dignity, respecting and promoting their autonomy; regulation, though often crafted with the purpose of treating individuals with dignity, often works to frustrate and limit individual autonomy.

If we ignore this task and instead argue on the substance of specific cases and policies, we will always lose the public case for sound policy. Sound policy may nonetheless emerge from serious and deliberative decision-makers, such as judges and those at the DOJ, FTC, and FCC. But so long as the common view favors regulation and distrusts the market, policy debates are a war of attrition.

The tenor of the Comcast/Time Warner Cable merger and Comcast/Netflix deal coverage has been a stark reminder that the discussions by those who understand these issues are driven by the concerns of those who don't. And it is no response at all to say that the average individual should know more. They shouldn't – they should know enough to trust the market in the abstract, so that they can focus their own efforts on developing specialties that will make the market work for others. The winning appeal to the public is that of dignity, of the market's power to protect and promote the interests of the individual – of all individuals – and of the counter-intuitive dangers that regulation poses to it.

Net neutrality: A complex issue, then and now

by: Gus Hurwitz

September 23, 2013 6:00 am

COMMUNICATIONS INTERNET

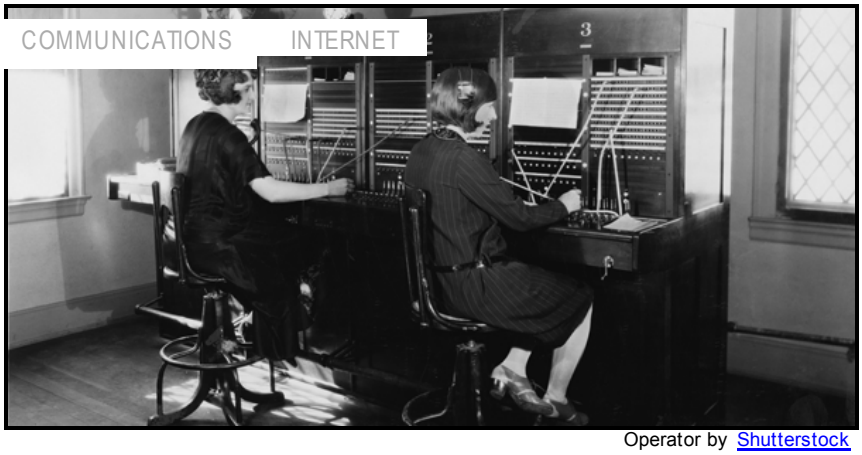

Operator by Shutterstock

tags: AT&T, FCC, net neutrality

Network Neutrality debates are fundamentally about switching – whether network switches can treat some packets differently from others. In this piece, I look back 100 years to the telephone interconnection debates of the early 20th century – and, in particular, to AT&T's preference for (non-neutral) manual switchboards over (neutral) automatic switches. This history reminds us that design decisions in complex networks are rarely as simple as network neutrality proponents suggest they are – and that market forces, if given time to operate, can secure the consumer benefits that regulators aspire to promote without the appurtenant risk that regulatory intervention may stunt the market.

The DC Circuit argument in Verizon's challenge to the FCC's Open Internet rules dominated tech and telecom news earlier this month. While analysis will surely continue – I'm unofficially obliged to nod to our own coverage here at TechPolicyDaily.com, as well as to the analysis of my colleagues at Truth on the Market (which has thus far proved eerily prescient, and where I am soon to post an extended discussion of the ideas discussed below) – I'd like to take this post to reflect on a ghost of telecommunications' past: the efforts – and fights – over interconnection of local telephone exchanges, and in particular the use of manual switchboard versus automatic switching, in the early 20th century.

Starting in 1907, interconnection was one of the great projects of the telephone industry. Up until that point, local exchange carriers frequently competed head to head in overlapping geographies. It was common for a single town to have both an AT&T-affiliated exchange and an independent exchange. And, amazing to modern readers, customers of one exchange could not call customers of the other!

AT&T changed its policies in 1907, in favor of what was then styled “universal service,” meaning that any telephone customer should be able to call any other telephone customer – an idea different than today’s meaning of the term. AT&T’s approach was largely to acquire competing local exchanges and to merge them into its own network or and to allow interconnection via its long-distance network. But it also agreed to, or was often subject to laws requiring, more direct interconnection between local exchanges (so-called “physical connection” of the exchanges).

Throughout this period, there were frequent complaints about AT&T operators discriminating against or dropping calls from or to customers on independent operators. These complaints were, of course, in addition to complaints that AT&T would not interconnect local exchanges (directly or via long distance lines), or charged too much to do so.

Adding to this, up until 1919 AT&T doggedly refused to adopt automatic switching technologies. Automatic switching (that is, mechanical switches, operated by the caller entering a number into his phone, and that therefore did not rely on switchboard operators) had been developed in 1888, patented in 1891, and widely adopted by independent exchanges over the next 20 years. Since automatic switches remove operators from the local exchange, they would have addressed many of the independents’ concerns about AT&T operators discriminating against their customers’ calls. But, until an operators’ strike in 1919 increased the labor cost of manual switchboards, AT&T refused to adopt this technology. Indeed, whenever it acquired an automatic independent exchange, it would replace that exchange’s automatic switch with a manual switchboard.

This history is an interesting precursor to the modern concept of network neutrality. AT&T, through its operators, allegedly offered preferential service to its own customers, blocked or dropped calls to or from its competitors’ networks, and offered a slower or otherwise lesser quality of service to calls to or from its competitors’ networks. (It is worth noting, too, that this was surely not the first case of non-neutrality-like conduct – for instance, the automatic switch was developed in response to alleged discrimination by a phone operator (directing business to her paramour), and the telegraph industry faced concerns about how it controlled access to news and information in the late 19th century.)

What does this tell us about the modern network neutrality discussion? The obvious answer is that it is an example of the monopolist service provider benefiting from blocking and discriminating against its competitors. Given that it could profit from disadvantaging its competitors, AT&T was happy to use a technology that facilitated such discrimination, and may even have encouraged it. This surely is the perception that AT&T’s competitors had.

But, then, as today, the situation was somewhat more complicated than the simple narrative suggests. While AT&T may not have minded that its switchboard operators were a burden to its competitors, it had independent reasons to operate in this way. The fact that AT&T would replace automatic switches with manual switchboards – even in exchanges where it had no competitors to disadvantage by such a move – demonstrates some other motivation.

There are several explanations for AT&T’s switchboard preference. The most common explanation is that

AT&T believed its customers preferred the experience of operator-mediated connections. While this may seem anachronistic today, it makes sense if you consider the disdain that many of us have for “interactive menus” – when you call up customer support, you almost certainly prefer to speak to a human over having to press a bunch of buttons.

But there were more reasons for AT&T to prefer manual switchboards. First, automatic switches couldn’t interface with its long distance service. Independent exchanges were, first and foremost, local exchanges that made some effort to offer basic long distance service, for who automatic switches’ inability to interface with long distance service presented little burden. ; AT&T was first and foremost a long distance carrier that interconnected local exchanges – its local exchanges had to seamlessly integrate into its long distance network. Related to this, its operators were trained, and developed an expertise, to place calls on its long distance network. They had a language, or API, that allowed them to establish long-distance calls quickly and efficiently. The independent exchange operators didn’t speak this language as well, so naturally received less efficient service. This language can be thought of both as an operating efficiency and a fixed cost.

AT&T was also investing heavily in building out and improving its long-distance service, and also continuing to build out its local exchanges. Developing and implementing a new exchange technology would have distracted from its long-distance efforts. And, much of its expansion, especially into rural areas, relied upon party-lines. These lines could not be billed using automatic switches.

A final reason for AT&T’s preference for manual switchboards was that operators were cheap. It wasn’t until the operators’ strike of 1919, in which 6,000 operators in the north east stopped work and ultimately negotiated higher salaries, that AT&T decided to adopt automatic switches, a decision that reflected not only the higher labor costs but also the fact that the operators had proved themselves to be potentially less reliable than the automatic switches.

This history, and these explanations, are surely not complete. Proponents of network neutrality can surely differentiate this history from the modern setting, much of AT&T’s conduct in this era is legitimately questionable, and the subsequent history (aided by technological change) vindicates many of the concerns about AT&T’s conduct (but also justifies much of AT&T’s conduct).

While this history doesn’t answer our modern questions about network neutrality, it does offer lessons. Perhaps most salient, we should remember that the networks of 100 years ago were far less complicated than today’s networks. So, too, were the services they offered. Yet even in that relatively simple setting the economics of those networks – and the factors that went into AT&T and the independents’ economic decisions – were quite a bit more complicated than usually understood. AT&T’s reluctance to embrace automatic switching resulted from its efforts to balance competing market incentives – to offer high-quality but low cost local and long distance service (best done with manual switchboards), on the one hand, with offering the better local but worse long distance service (with automatic switches) on the other. These are exactly the issues we want firms to struggle with.

Network neutrality presents much the same question: how to best provide switching to different types of

services. We've been here before. As in the 1910s, the concern is on the one hand that broadband providers' switching decisions can harm other firms but on the other hand that they can better facilitate high-value services with particular network requirements. These tradeoffs exist today, and will exist 100 years from now – they are inherent in any complex network.

One hundred years ago, AT&T struggled with these issues and ultimately adopted automatic switches – it did so in response to market forces, not regulatory fiat. It is scary to think how things may have been different had the government intervened and specified which switching technologies AT&T had to use — free from such government oversight, Bell Labs went on to develop the first electronic switches, which became the first computers, and helped pave the way for the switches that power the Internet today.

Yet today, despite the market being more competitive and subject to more scrutiny than AT&T was 100 years ago, network neutrality proponents advocate regulatory control of switching – government edicts saying what switching technologies can and cannot be used (and therefore, can or cannot be further developed). And this despite little evidence that network neutrality poses anything more than hypothetical harms. The market has worked in the past – how about we give it a chance today, before deciding to regulate it? If things don't work out, if the net-neutrality proponents' parade of horribles does come to pass, regulatory intervention will still be an option. Until then, we should be cautious of the impulsive regulatory instinct: too-readily opting for regulation over the market is the sort of automatic switch we should avoid.

Return to the Article

January 16, 2014

The Net Neutrality Decision Should Have Both Sides Happy

By Gus Hurwitz

This week's decision from the DC Circuit Court of Appeals, which rejected the Open Internet Order's no-blocking and non-discrimination rules, is very important. But, despite its importance, it is really not all that surprising. The more interesting question is what comes next. FCC Chairman Wheeler appears to have shown his hand, suggesting strongly that he prefers the Commission use its newly affirmed Section 706 authority to bring adjudicatory enforcement actions against firms that may engage in possibly problematic conduct, instead of continuing in the Commission's efforts to develop prophylactic network neutrality rules. This is the correct approach for the Commission to take. Let's look at why.

In its decision, the DC Circuit applied standard principles of administrative law (the area of law that governs the division of power between courts and agencies). These principles give agencies substantial power to interpret their statutes broadly. They, however, do not give agencies power to contravene clear statutory requirements. Under these principles, which were strongly reaffirmed only months before Oral Arguments in the Supreme Court's City of Arlington opinion, it is not surprising that the DC Circuit found that the Commission has sufficiently broad authority under Section 706 to regulate aspects of the Internet; nor is it surprising that the DC Circuit found that portions of the Open Internet order violated the Communications Act's clear prohibition on imposing common carriage requirements on networks like the Internet (which the FCC classifies as an Information Service). Indeed, this is precisely what most commentators have expected the judges would do.

I don't mean to overstate my, or anyone's, powers of prediction. There was still a lot of play in the case's likely outcome: would the judges find the no-blocking rule ran afoul of the common carriage restriction, or only the non-discrimination rule? Would all three judges sign on to the majority decision, or only two? Would the decision say anything substantive about Verizon's First Amendment or Takings arguments? And, of course, there was always a possibility that the judges could do something entirely unexpected. But, by and large, it was widely expected that the DC Circuit would find the FCC has broad authority to regulate the Internet under Section 706 but that it was statutorily barred from using that authority in some of the specific ways that it had.

The interesting question is "what next?" FCC Chairman Tom Wheeler may have already answered this question for us. In a posting on the FCC's blog, the Chairman explains that his "strong preference is to [use the Commission's power] in a common law fashion, taking account of and learning from the particular facts that have given rise to concern. ... If something appears to go wrong in a material, not a trivial, way, the FCC will be available to use the totality of its authority for adjudication and enforcement." This strongly suggests - borders on stating outright - that the Chairman does not intend to craft new, prophylactic, Network Neutrality rules. Rather, he strongly prefers a wait-and-see approach, allowing firms to develop new business models and practices. Should any of these raise the ire of the FCC, the Commission will use its (clearly articulated) Section 706 power to reign in the problematic conduct.

This is exactly the approach that the FCC should - and always should have - taken. A fundamental problem for Network Neutrality proponents is that the harms over which they are concerned are speculative. There have only been a few arguable instances of firms engaging in potentially problematic conduct - consumer harm in these cases has been de minimis, and their resolution has not required network neutrality rules. On the other hand, there is substantial evidence that some business models and practices prohibited by the Open Internet order could yield substantial benefits to consumers. Rather than proscribe entire classes of potentially beneficial conduct out of speculative concerns of possible harm, the FCC should take the more reasoned and cautious approach of waiting to see how the market actually develops. After yesterday's decision, it is clear that the Commission has the authority to act if action proves to be needed.

The most challenging aspects of how the Commission should proceed stem from administrative law, not communications law. As a bit of an aside, it is useful to remember that the DC Circuit's decision was primarily based in administrative, not communications, law. Indeed, the concerns that Judge Silberman express in his dissenting opinion are very apt: the Communications Act gives the FCC alarmingly broad powers, broader, perhaps than any other agency has ever been granted by Congress. It is well past time for Congress to review - and

refocus - these powers. While I agree with the majority's holding that the Commission's interpretation of its Section 706 powers is permissible, this does not mean that I agree that it is proper for the Commission to have been delegated such broad powers.

Nevertheless, if the Commission is to use these powers well, it should undertake a process to more precisely define the contours of its Section 706 powers. This is essential both for industry to know what conduct may or may not be problematic, and also for the FCC to take action against conduct that it deems problematic. For instance, and in particular, any future FCC enforcement action likely turns on what it means that the FCC shall "encourage the deployment ... of advanced telecommunications capability to all Americans ... by utilizing ... measures that promote competition in the local telecommunications market." If the FCC has not engaged in serious and public discussion - something substantial but that may fall short of a rulemaking - about what this means prior to its decision to take action, it may well be difficult for the Commission to prevail in any enforcement action. Just as administrative law gives the Commission authority to construe Section 706 broadly, it also denies the Commission the ability to construe Section 706 favorably for the sake of litigation and requires the Commission to give firms sufficient notice of how it will construe Section 706 prior to taking any enforcement action.

The DC Circuit's decision, particularly taken with Chairman Wheeler's likely preference for an ex post adjudicatory approach to future concerns, should leave both Net Neutrality proponents and opponents reasonably happy. The Decision affirms that the Commission does have authority to take action against firms engaging in harmful conduct. And given the consumer benefits likely to result from conduct previously proscribed by the Open Internet Order, the Commission is unlikely to take action against firms using their newfound freedom to develop innovative and pro-consumer business models. The breadth and structure of the Communications Act is still problematic in many ways - this is a problem for Congress to address - but administrative law's procedural safeguards should provide some measure of stability and comfort for the firms that have previously been hampered by the Open Internet Order.

checkTextResizerCookie('article_body');

Gus Hurwitz is a visiting fellow at the Center for Internet, Communications and Technology Policy at the American Enterprise Institute, and an assistant professor at the University of Nebraska College of Law. He is a contributor to TechPolicyDaily.com.

Should we regulate firms we just don't like?

People don't like telecommunications companies much. Telephone, cable, cellular, or Internet – these are companies that people love to hate. Curiously, people dislike these companies despite loving the services that they provide. And they are disliked despite the substantial role they play and have played in developing the modern economy – both as innovators and platforms for innovation.

I spend a lot of time trying to understand why people so dislike these firms. There are many possible reasons – some good, most bad. The thing that I find most remarkable about our collective dislike for these firms is how innate and visceral it is. Importantly – and this is the idea that I want to take a moment to develop – this dislike is not new. It is illustrative to look back to the pre-telephone era: in the era of the telegraph, people didn't much like Western Union, the telegraph monopoly, and their dislike looked much then like the dislike of telecommunications looks today.

I don't mean to say that these firms are angels. To the contrary, their heritage is very often that of state sanctioned or regulated monopoly. With a guaranteed customer base and ensured revenues, such firms often focus more on technology and networks than on customers and service – and sometimes they focus on none of the above. That's the nature of the incentives created by state-sanctioned and regulatory monopolies, and the lesson of decades of efforts to work within the regulated monopoly model. Even in today's increasingly competitive telecommunications environment, many firms' approach to the customer is rooted in a culture developed during the era of regulated monopoly – and it is too often the case that the regulated monopoly remains the status quo, still protected by state and local franchising authorities.

On the other hand, much of the antipathy towards these firms is ill-placed. In high fixed-cost/ low marginal-cost industries such as these, prices are almost always above marginal cost, such that users almost always feel that they are being over-charged. What's more, the marginal costs are often so low that they feel like they are zero – users feel like they're being charged too much for something that should be free. Additionally, the grass almost always seems like it is greener on other networks. Given the pace of innovation, the depreciation cycle, and the time it takes to upgrade networks, there will almost always be other networks available that are better, cheaper, or both. Users on any given network can almost always point to another network as an example that they are paying more for less than folks on other networks. And add to this that users rarely care about the network: they care about the content that it provides. A happy user rewards the content delivered by the network; an unhappy user blames the network for its failure to deliver desirable content. Once again, it's a lose-lose world for telecommunications firms. And, because these industries are relatively concentrated (either, or both, because they are characterized by substantial economies of scale and because of state licensing requirements that keep competitors out), it is easy for consumes to point to these firms as monopolies, blaming these problems on their size, and not on the nature of the services they offer.

All of these reasons for users to dislike telecommunications firms are premised on their actually being a customer of those firms – they dislike the firms because they are, or at least believe they are, receiving sub-par service at too-high prices.

Let us compare this to how people viewed Western Union during the era of the telegraph. First, we must note that telegraphs were not a widely used service. Sending a telegraph was relatively expensive and inconvenient, generally requiring the sender to go to the local Western Union office. Unlike modern forms of telecommunications, the telegraph was not something used by ordinary people

on a day-to-day basis.

This makes popular views of Western Union very interesting. They reflect abstract notions about these firms, independent of any actual conduct or harm they may cause. As explained in by Menahem Blondeim in an article exploring Congressional efforts to regulate Western Union and the Associated Press in the late 19th century, at that time monopolists like Western Union were viewed as "mutations, the consequence of some sinister tampering with the natural order of things. [They were] not merely economic entities but powerful new political forces which must be opposed in the name of American democracy." Menahem Blondheim, *Rehearsal for Media Regulation: Congress Versus the Telegraph-News Monopoly, 1866-1900*, 56 Federal Communications Law Journal 299 (2004) (quoting business historian Thomas McCraw).

Yet, despite this dislike of monopolies, Congress, after decades of effort, was never able to find a way to regulate Western Union. Between 1866 and 1900, Congress considered 96 bills and resolutions to regulate Western Union; none passed. During all this time Congress was unable to get beyond the mere feeling that Western Union was the result of a sinister tampering with the natural order of things to "articulate a clear rationale for why something needed doing. [W]ould-be regulators could not sustain the allegation that Western Union was leveraging its monopoly position to charge unfair prices for its services. Additionally, they could not demonstrate that the corporation retarded technological development in order to buttress its monopoly. Even the most populist voices that joined in the debate could not convincingly point to a pattern of corrupt or even unfair business practices of the corporation."

To complete this snippet of history, over time Congress's attention was drawn to the nature of Western Union's business: as a communications monopoly, Western Union "had the power to shape what the people knew, [and] was presumed to wield power over what they thought. While general, abstract notions of [monopolies] as dangerous aberrations hardly provided a mandate for legislative action, a monopoly of knowledge did." This, in turn, led legislators and the courts to focus on the relationship between Western Union and the Associated Press, and in particular on the Associated Press and its monopoly over the news – a monopoly that was ultimately broken up in 1900. By that time attention was shifting from telegraphy to telephony and Congress's interest in regulating Western Union waned.

The lesson to take from this is that we have a long history of viewing telecommunications firms through jaundiced eyes. Even where users have no direct contact with these firms, the fact of their size paints them as sinister mutations – even where they have had no opportunity to cause any harm, their malfeasance is assumed. This is compounded by what they carry: information, knowledge, the fuel of democracy. Having assumed their malfeasance, we now find the harm proved, as well.

It is entirely possible for these firms to, in fact, engage in anticompetitive, or otherwise harmful conduct. This often follows from the economic structure of the telecommunications industry; and it often follows, or is exacerbated, by governmental intervention in creating state sanctioned or regulated monopolies. But given our track record of innate and visceral dislike for these firms, we must be sure to judge them fairly.

Lessons from behavioral economics should apply here: where we are concerned about the conduct of telecommunications firms, let us recognize our propensity to demonize them. We should demand actual evidence of actual harm, and not accept mere speculation and innuendo.